SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAG

AD-A227 364

DTIC ELECTE SEP 26 1990 S D CS

1a. REPORT SECURITY CLASSIFICATION	1b. RES
2a. SECURITY CLASSIFICATION AUTHORITY	3. DIST Unlimited distribution Approved for Public release
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE	
4. PERFORMING ORGANIZATION REPORT NUMBER(S) Grant AFOSR-89-0072	5. MONITORING ORGANIZATION REPORT NUMBER(S) AFOSR-TR- 90 0944

6a. NAME OF PERFORMING ORGANIZATION University of California, San Diego	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION AFOSR
6c. ADDRESS (City, State, and ZIP Code) Mail code S-034 La Jolla, CA 92093		7b. ADDRESS (City, State, and ZIP Code) Bldg 410 Bolling AFB, DC 20332-6448
8a. NAME OF FUNDING / SPONSORING ORGANIZATION Air Force/NM	8b. OFFICE SYMBOL (If applicable) NM	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER AFOSR - 89-0072
8c. ADDRESS (City, State, and ZIP Code) Building 410 Bolling AFB DC 20332-6448		10. SOURCE OF FUNDING NUMBERS

PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT ACCESSION NO.
61102F	2304/K7	A4	

11. TITLE (Include Security Classification)
Investigations of Equilibria, Lattices, and Chatoic Dynamics of 2-D Hamiltonian Point Vortices

12. PERSONAL AUTHOR(S)
Dr. James Kadtke

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Year, Month, Day)	15. PAGE COUNT
Final Technical	FROM 11/1/88 TO 10/31/89	1990, August, 8	

16. SUPPLEMENTARY NOTATION

17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

The final report summarizes the work performed under the AFOSR grant number AFOSR-89-0072, originally entitled, "Equilibria, Lattices, and Chaotic Dynamics of Point Vortices." The aim of the proposed research effort involved theoretical work, both analytic and numerical, on a number of different problems which were all loosely tied together as involving some aspect of vortex systems, and their relation to chaos in fluid flows. Significant results were obtained during this funding period in several major topics. The first topic which was investigated was a continuation of the author's previous work on vortex lattices. Results consisted of the refinement of the analytic expression for the lattice summation of an infinite lattice of point vortices, and use of this expression to calculate the allowed lattice structures of the two-component triangular lattice. It was also shown how these expressions can be used to calculate the bulk physical properties of vortex lattices, by calculating the energy of slip displacement for the triangular lattice.

20. DISTRIBUTION / AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT. ☐ DTIC USERS	21. ABSTRACT SECURITY CLASSIFICATION Unclassified
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. A. Nachman	22b. TELEPHONE (Include Area Code) (202) 767-4939 / 22c. OFFICE SYMBOL NM

D Form 1473, JUN 86 — Previous editions are obsolete. — SECURITY CLASSIFICATION OF THIS PAGE

FINAL REPORT:

AFOSR Grant AFOSR-89-0072

PI: J.B.Kadtke, Ph.D.

Accesion For
NTIS CRA&I
DTIC TAB
Unannounced
Justification
By
Distribution /
Availability Codes
Dist
Avail and / or
Special
A-1

Best
Available
Copy

Table of Contents

I. Introduction

This final report summarizes the work performed under the AFOSR grant number AFOSR-89-0072, originally entitled "Equilibria, Lattices, and Chaotic Dynamics of Point Vortices". The aim of the proposed research effort involved theoretical work, both analytic and numerical, on a number of different problems which were all loosely tied together as involving some aspect of vortex systems, and their relation to chaos in fluid flows. Significant results were obtained during this funding period in several major topics. The first topic which was investigated was a continuation of the authors previous work on vortex lattices. Results consisted of the refinement of the analytic expression for the lattice summation of an *infinite* lattice of point vortices, and use of this expression to calculate the allowed lattice structures of the two-component triangular lattice. It was also shown how these expressions can be used to calculate the bulk physical properties of vortex lattices, by calculating the energy of slip displacement for the triangular lattice.

A second major topic which was researched during this period was the investigation of chaotic motion in fluid flows due to vortex dynamics. One aspect which was investigated was to demonstrate the presence of low dimensional chaos in an actual experimental open flow. This work used experimental data from the erratic fluid flow downstream of a cylinder, obtained from INLS's fluids lab, and utilized a new technique developed by the author for measuring the dimension of a chaotic system from a time series. The motion was found to have an apparent dimension of four, which agrees with an alternative model of the system. Another problem involving chaos due to vortices was an analytic and numerical investigation of the motion resulting from the interaction of a vortex in an open flow and a stationary bluff body, with small sinusoidal perturbation. It was found that, for the proper parameter regimes, the passing vortex could actually be 'chaotically trapped' around the body for significant periods of time. This behaviour is a relatively new phenomenon and one of possibly large significance.

Finally, motivated by an interest in measuring and describing low-dimensional chaos in vortex and fluid flows in general, a significant amount results were obtained in the general theory of nonlinear analysis. An extensive method was developed for prediction of chaotic motion based on a global, functional description of the attractor, and utilizing only a time series of data as

input. This method was found to have superior predictive ability over most commonly used methods. Another significant result was the development of a new method for the determination of the minimum embedding dimension necessary to reconstruct the attractor for a system. This method is an entirely new approach based on information theory, and offers an alternative technique to the ubiquitous Grassberger-Procaccia algorithm (its initial uses were to analyze the experimental data mentioned above). In addition to these two major projects, results were also obtained in work done on the investigation of a new type of dynamical mapping, which has interesting property of being locally conservative but globally dissipative. Lastly, work was begun and is still ongoing in developing a new method for calculating the lyapunov exponents of a chaotic system from a time series of data, in the presence of additive guassian noise.

In all, work performed during the funding period has resulted in three published papers, two papers currently in review, and two papers in preparation (all of these are included as appendices). Versions of the pre-prints and published papers will be supplied as they become available.

II. Vortex Lattices

This section describes analytic and numerical work done in the calculation of the lattice summations for the energy of a two-dimensional, infinite lattice of point vortices. This work was done in collaboration with L.J.Campbell and M.M.Doria of Los Alamos National Laboratory, and Physical Review publication resulting from this work appears as Appendix A.

The calculation of properties (such as the energy density) for infinite lattices with Coulomb interactions has a long history. In particular, the two-dimensional case which arises from lattices of point vortices is an example of this problem, as both satisfy a logarithmic potential, and with the only difference being that of the resulting dynamics. Vortex lattices are of importance both because of their mathematical properties, as well as being reasonable models for superfluid helium systems, systems of line charges or currents, screw dislocations in crystals, vortex-like interactions in the quantum hall effect, and most recently as a mechanism for 'high Tc' superconductors. Until the present work, no closed form expression existed for these lattice summations, although there had been several attempts to derive them(see references in Appendix A). Numerical simulations thus consisted only of clever ways to do the summations explicitly.

The principle result of the work discussed here is the development of an expression for the energy density of a lattice of point vortices in terms of a very rapidly convergent product expansion. Although a more primitive form of the summation was developed prior to this funding period, these results represent a considerable improvement in the formalism involved, as well a a more sound understanding of the physical interpretation. In particular, the correct form of the normalization was finally understood, as well as a more general understanding of the form and interpretation of the artificial neutralizing background one must add to cancel the mathematical singularities involved in the infinite summation (see the introduction in Appendix A). The final results of this reformulation is summarized in Eq.21 of Appendix A; this equation gives the *energy density* of of a lattice of vortices with given species of vortices of given strengths. This equation allows for any arbitrarily shaped, four-sided unit cell and also for arbitrary numbers and strengths of vortices. The correct normalization now allows for the correct comparison between different lattices with similar number densities for different species. Since derivatives and other

operations can easily be performed on the expression, it is also suitable for the calculation of bulk properties of the lattices and even to investigate lattice dynamics. In its present form this expression is directly applicable to a wide variety of problems associated with this type of logarithmic-potential lattice, such as those mentioned above.

In addition to the above results, this formulation was then used to obtain several new results for vortex lattices. Previously, only properties of the single-species, single-vortex square and triangular lattices could be calculated. Section IV of Appendix A presents the lattice shapes allowable for several new varieties of vortex lattices. These shapes are all calculated by similar means: minima of the lattice energy density are found by sweeping through the lattice parameters, and actual lattice configurations are assumed to exist for these minima. Section IV shows several new lattice structures, and gives some of the lattice energies, for which the energy of the square and triangular lattices is verified. It should be noted that many other new configurations of various types of lattices have been generated, and that only the two dimensional case has been presented in the paper. Other configurations may eventually be published elsewhere.

Finally, some results were obtained on the calculation of some bulk lattice properties using the energy density formulation. The property of significant physical interest for the lattices of this particular type was the energy of slip displacement. This energy in effect measures the 'rigidness' of the lattice to deformation along one of its principal directions. Figure 3 of Appendix A presents a summary of numerical calculations done to measure the energy of slip displacement for a triangular vortex lattice. Calculations of this type were entirely impossible previously. This calculation is also primarily useful as an example of the practicality of Eq.21 for the calculation of a wide variety of properties associated with logarithmic-potential lattices.

It should be noted that although the author has not pursued work on this project much beyond what is presented here, L.J.Campbell has demonstrated the importance of these results by continuing to apply this formalism to a number of outstanding problems, in particular some aspects relating to models of high Tc superconductors.

III. Vortex Chaos

A. 'Chaotic Trapping' in Open Flows

This section describes work in progress which has been ongoing for approximately one year, and which has been done in collaboration with E.A.Novikov of INLS, UCSD. The original analytic foundation for the work appears in a short pre-print by Novikov, and in this continuation of the work this author has completed an extensive numerical study which will shortly appear in pre-print form, and which contains new and significant results concerning the interaction of vortices with bluff bodies in open flows. Examples of the numerical results appear in Appendix B, as well as Novikov's original pre-print to include background discussion (some of these numerical results were also recently presented in a talk at SIAM's Dynamical Systems conference in Orlando, May 1990).

The underlying model for this investigation is that of a bluff body (in this case a two-dimensional cylinder) in a uniform two-dimensional flow, with a single vortex passing by and interacting with the cylinder. To this system is added a small perturbation which consists of sinusoidal vibration of the body along the flow direction. This system is a simple and general model for a vortex-like structure interacting with a body in an open flow, for which there are many physical analogous; the most obvious of these is that of a tornado interacting with structures such as large buildings. In the case of zero perturbation, the vortex trajectories, and hence the topology of the flow field, can be solved exactly (see Appendix B). However, for finite perturbation and the proper parameter regimes, Novikov has shown analytically that one can generically expect chaotic motion of the vortex to emerge. This chaotic motion results in several consequences, the first and most obvious of which is that the resulting vortex trajectory is unpredictable and extremely sensitive to initial conditions. Hence the direction of scatter of the impinging vortex is highly erratic with respect to initila position. A consequence of this is another important phenomenon, namely that of 'chaotic trapping' (a related phenomenon has been introduced previously in gravitational interactions by M.Henon). This phenomenon consists of vortex motion for which the vortex approaching from infinity can become trapped in rotational motion around the body for a (finite) period of time, and then escape again. The trapping is a result of the vortex

being caught in the stochastic layers around the flow separatrix, which is generated by the perturbation. Because this resulting motion is chaotic, the rotational motion of the trapped vortex can in turn result in large pressure differences on the boundary of the body. This phenomenon is apparently new, and because the model predicts that this is the generic state for finite perturbation of such systems, this phenomenon should be recognizable in physical systems as well. Speculation is that this behaviour may be related to the destructive ability of tornadoes, as well as that of the highly dangerous downdrafts which sometimes occur around airports. In addition to these physical applications, the system is also of considerable interest as perhaps the simplest model of an open flow system which exhibits such rich chaotic behaviour.

The numerical investigations of the above system have produced a large body of results which indicate an even richer behaviour than was suggested by the analytic analysis. The first major result was the study of the variation of the topology of the flow field, with zero perturbation, as the dimensionless parameter σ of the dynamical equation is changed (see paragraph 5 of the Novikov preprint). Roughly, this parameter measures the ratio of the vortex and flow field strengths. To study the variation of topology, the positions of the stable and unstable stagnation points of the flow were found for the entire range of σ, using computer algebra and solving numerically for the roots. Once distinct regimes for the root positions were identified, the flow fields were mapped out by integrating the dynamical equations. Using this method, twelve distinct topologies were identified for the unperturbed flow, far richer than was first suspected. These topologies are outlined and included in Appendix B. The majority of these topologies are capable of exhibiting stochastic regions, and hence trapping phenomena. In addition, Case 3 of these topologies indicates a stagnation point whose character is a mixture of hyperbolic and elliptic, which may itself be a new type of structure in Hamiltonian flows.

The second important result of the numerical investigations was the verification of the existence of the chaotic-trapping phenomenon. Since the size of the stochastic region, and hence the probability of trapping, is moderately dependent upon the frequency of the perturbation, the Melnikov integral for the system was computed numerically to determine the frequency range corresponding to large stochastic regions. Using different values within this parameter regime, a large number of trajectories were found, by numerical integration, which exhibited trapping behaviour. Two typical examples are

shown in Appendix B. Trajectories were found which became trapped, performed as many as twenty erratic revolutions around the body, and then escaped. Within the proper parameter regimes, the measure of initial conditions resulting in trapped trajectories seems to be quite significant, as these trajectories were relatively easily found.

Related to the above phenomenon, a third interesting result was obtained for this system. For the case where a stable and unstable stagnation point lie somewhat near the boundary of the body and on the same side, trajectories were found where the vortex could actually first become trapped around the body, then around the elliptic point, and then often switch back and forth several times. These 'switching' trajectories have been observed for several initial conditions, and seem to be a somewhat unusual and counter-intuitive result. An example of such a switching trajectory is also shown in Appendix B.

Finally, for a few cases of different parameter values, Poincare sections of the vortex motion for specific trapping trajectories were taken. This was done solely to aid in gaining intuition about the structure and appearence of the stochastic layers which cause the trapping phenomenon itself. An example of one of these sections is included in Appendix B. Although not of direct relevance to the analysis discussed above, this type of chaotic analysis will be the subject of a future more detailed investigation.

Future work on this project is planned to be quite extensive. Several numerical experiments are planned to characterize the stochastic nature of the system, including determining the measure of initial conditions which result in trapping trajectories (it is suspected that this may result in a 'devil's staircase'), measuring time series of boundary pressure on the body for trapping trajectories, and further Poincare analysis. Beyond that, the system will be generalized by considering vortex-dipole interactions with the body, more general boundary geometries, and alterations of the model to make it more applicable to specific physical systems.

B. Low-Dimensional Chaos in an Open Flow Experiment

In the original proposal for this work, it was mentioned that the possibility existed of developing new results for the generalized von Karman street as a better model of the wake of an open flow past a bluff body. Although new analytic results could not be obtained, an investigation was done of some experimental data of such a flow, which seemed to exhibit a chaotic nature. The principle tool for the analysis was a new technique for determining the minimum embedding dimension of a chaotic signal, developed in part by the author and which is separately presented in this document in Section IV. Using this method, the region of the flow investigated was found to be chaotic with a dimensionality of four, which also agrees with a first order model proposed for the system. A summary of these results are given on page eight of the pre-print 'Information Theoretic Methods for Determining Minimum Embedding Dimensions of Strange Attractors', which is included in Appendix E. As with the previous section, this work is still in progress, and a more extensive and cooperative effort is planned to identify low-dimensional chaos in various aspects of these flows.

The experiment from which the data was obtained was performed by M.Gharib and K.Lewis at the DARPA/URI fluid dynamics laboratory at INLS, UCSD. A detailed description of the experiment will appear as a pre-print in the near future. Briefly, two thin rotating cylinders of slightly different radius are placed end-to-end and perpendicular to an otherwise uniform flow field. Both cylinders generate vortices in their downstream wake, however the mismatch in radii causes an unstable interaction which results in low-dimensional chaos in the flow near the boundary of the two regions. A photograph of the flow, supplied by K.Lewis, for a typical experimental run is shown in Appendix C, as are a time series and FFT for the chaotic region.

The MDL technique for measuring minimum embedding dimension was used on a set of several different time series from these experiments, in an attempt to determine whether this flow was indeed low-dimensional chaos (this method is described more fully in Section IV). It should be noted here that the conventional method for determining embedding dimension, ie. the Grassberger-Procaccia algorithm, yielded inconsistent results for this particular system, and this was one of the primary motivations for the use of the new MDL algorithm

to analyze this data. After considerable analysis, it was shown that this flow typically seemed tc have a dimensionality of four, although the dimension could be as high as six for some parameter regimes. A plot of the MDL function for a typical data series is also shown in Appendix C, showing the minimum of the MDL function occuring at a dimension of four. Since Gharib and Lewis have proposed a model for the chaotic region based on coupled duffing oscillators, which should also have an expected diemnsionality of four, these results seemed to confirm this conclusion.

Although work on this project is still ongoing and results are somewhat preliminary, this investigation could prove very significant as an excellent example of low-dimensional chaos in a flow which can be demonstrated experimentally, analytically, and through nonlinear time series analysis. Future work is also planned, in cooperation with researchers in the DARPA fluids lab, which will involve a similar analysis in an attempt to look for low-dimensional chaos in the velocity and acoustic fields of a submerged jet.

IV. Measurement and Prediction in Chaotic Systems

Because of the authors interest in investigating chaos in vortex flows and in fluid flows in general, there was considerable motivation during this project to study the various methods for measuring, characterizing, and predicting chaotic behaviour in flows, in particular those resulting in time series from actual physical systems. It was generally found that, although idealized methods for analyzing chaotic motions were well established, application of these techniques to real physical data, which often include significant levels of noise, was often very poor. A significant amount of effort during this funding period was therefore devoted to investigating and developing improvements in these techniques, which would be more robust to the problems associated with real data (ie. noise, short data sets, irregular sampling of data, etc.). These new techniques were investigated with an eye towards increasing the analytic abilities for experimental data, especially that being generated at the DARPA UCSD fluids laboratory. This work during the past 1 1/2 years has resulted in two new and powerful methods for the analysis of chaotic data from actual physical systems, and has produced two published papers and one paper currently in review. These two methods are outlined below.

A. Global Prediction for Dissipative, Chaotic Systems

Most methods of analysis for chaotic systems involve the measurement of several physical properties which are known to indicate chaotic motion, such as the dimension, the lyapunov spectrum, etc. Once these quantities are determined, the obvious next step should be to use these quantifiers for model selection, signal processing, prediction, and the like. Surprisingly, only a handful of papers exist which attempt to develop any of these applications, and very few serious efforts to analyze actual data have been done. The work described below, done in collaboration with H.D.I.Abarbanel and R.Brown of INLS, UCSD, was an attempt to develop a general method for modelling a chaotic systems' flow in phase space by simultaneously utilizing as much dynamical information as possible, in the most efficient way, and then utilizing this information for prediction. Prediction in this sense means the generalized forward extrapolation of short segments of phase space trajectories, which can then be used for actual prediction or for signal processing, noise reduction, etc.

The results of this investigation grew into a somewhat extensive general methodology, which is described in the two papers included here as Appendix D.

To give a brief overview of the method, the basic idea starts by reconstructing a chaotic systems' attractor in phase space from a time series using time-delay embedding and the standard methods to determine the embedding dimension and the autocorrelation time. The general procedure is to model the local flow on the attractor by a mapping function which uses the information of where nearby neighboring points on the attractor are mapped to. This function can then be used to predict where a new point on the attractor will evolve. In the formulation and the numerical algorithms which were developed, a general form of the mapping function is used with the important properties that it is global, in the sense that only one function is used over the entire attractor, and secondly that the function is well defined analytically, so that for example gradients can be computed and local derivatives can be used for parameter fitting. Specific information about the given dynamical system is then built into the mapping function in the following way: the mapping function contains parameters which weight the way local information on the attractor is used, such as the number of neighbors to include, the number of previous iterates to include, the length scale involved, etc. The correct parameters that globally describe the particular attractor are then found by nonlinear least-squares fitting of the function to the data set taken from the dynamical system. The correct parameters are then chosen as those which optimally reproduce the given data set from the system.

Perhaps the most important aspect of this method is the inclusion of additional dynamical information about the system in the mapping function. In addition to the dimension, there also exist other important quantifiers of the character of the chaotic motion which can be extracted from the time series data. These include the lyapunov spectrum of exponents, and the probability density of orbits on the attractor. These quantities are measured from the time series using standard algorithms, and are then also built into the mapping function in the following manner: the functional values that these quantifiers should yield for the given mapping are derived from it by taking gradients and using appropriate definitions. When the nonlinear least-squares parameter fitting is being done, the map is simultaneously constrained to also reproduce the appropriate values of the dynamical quantifiers. Examples of the results of this type of constrained optimization are given in both papers in Appendix D. The point of this is that

the resulting map now reproduces not only the data to some accuracy, but also reproduces the correct dynamical characteristics of the system. It is therefore expected that this constrained fit for the function will prove to be a more accurate predictor of the actual motion of the system.

Extensive testing and development of this idea have generally shown that this type of constrained parameter fitting seems to produce considerably superior predictive power than that of other less sophisticated techniques. Results of this work have generally shown (see Appendix D) that it is possible to find global mapping functions which reproduce time series data to an excellent degree (0.5% average rms error or better) and which are also capable of accurately reproducing the dynamical quantifiers, and hence have the same chaotic invariants, as the original system. Additionally, numerical experiments with the reproduction of known trajectories have shown that this type of predictor can accurately predict orbits significantly farther than the majority of other known methods, with shadowing trajectories for example staying close to the original Henon trajectories for as long as seven or eight iterations. In addition, this method seems better suited than most methods to actual experimental data, which may include significant noise components, because of the inherent averaging of the local phase space flow which occurs via the mapping function.

Although much of the work on this technique is completed, it has yet to see widespread use for time series of actual physical systems. Therefore, future plans for this project involve primarily identifying systems for which this technique may prove useful, and for which more practical experience can be gained. Actual improvements to the method will involve investigations of different types of general mapping functions, which could improve accuracy, and also of different algorithms for searching data sets for nearest neighbors, which currently consumes the majority of the computational resources associated with the calculations.

B. Information Theoretic Methods for Determining Dimension

The second major pro[illegible]hich was undertaken to develop new methods for chaotic time series analysis was the attempt to develop a method for determining the embedding dimension of a chaotic attractor. Determining the correct embedding dimension is the first necessary prerequisite for performing time-delay reconstruction of the attractor from a time series. The standard, most successful method for determining the embedding dimension is the Grassberger-Procaccia algorithm, but this algorithm has several well known difficulties, including ambiguity in determining the *minimum* embedding dimension that can be used, as well as sensitivity to noise, and often quite substantial data requirements. The new method, which was developed in collaboration with H.D.I.Abarbanel of INLS, is based on a result from information theory, and has the advantages of unambiguously determining the minimum embedding dimension allowable, as well as being much more robust to noise and requiring less data. The background and results of this new method are presented in the pre-print included in Appendix E, which is currently in review.

The central tool of the new method is a result from information theory which was developed over some time by Aikaike, Wax, Kailath, and others. This result is the definition of a function, called the Minimum Description Length (MDL) function, which quantitatively weighs a functional which is essentially the maximum likelihood fit to a data set, versus a measure of the complexity of the model used to generate the fit. In simpler terms, it weighs the trade-off between a data model being a better fit to the data, versus how complex (ie. how many parameters) the model has. For a class of fitting functions whose dimensionality is a variable, the MDL function can be proven to take a minimum at the minimal number of dimensions necessary to describe the data. The principal result of the paper in Appendix E is the adaptation of this function for time series of real data and for different normalizations and parameter counting, and is now called the Data Description Length (DDL) function. This paper demonstrates how the DDL function can be used to determine unambiguously the minimum dimension necessary to embed an attractor from a time series of data.

The actual algorithm which was developed works by first constructing a data matrix from the time series, and then calculating its eigenvalues by singular value decomposition. These eigenvalues are used in the determination of the

maximum likelihood fit to the data. For the present method, a guassian distribution of the reconstructed attractor is assumed, which of course is inaccurate for most attractors. The central point here is that although our assumption is crude, the question being asked is also rough, ie. how many independent dimensions is the attractor distributed along. By assuming a guassian distribution, the specific form of the DDL function can be written down, and from this formula the numerical value of the DDL function can be calculated for each value of the embedding dimension. The dimension for which the DDL function takes a minimum is the dimension which best describes the data, under the constraints of the particular distribution chosen. The algorithm has been developed to the point where one simply supplies the time series as data, adjusts a few parameters to the proper regimes, and then picks off the minimum of the resulting plot of the DDL function to find the correct embedding dimension.

Results of experimentation with the DDL algorithm (see Appendix E) have shown the technique in most cases to unambiguously yield the correct embedding dimension for test chaotic systems. In addition, and quite importantly, the method works in the presence of significant amounts of noise (up to 15-20 dB SNR), and also requires far less data and computational resources than the Grassberger-Procaccia algorithm. This makes the method particularly well suited to analyzing experimental time series. There are still some difficulties, however, with determining the dimension of attractors whose topology is not simple (ie. multiple lobes or interleaving). This problem is almost certainly related to the choice of a gaussian distribution for the underlying maximum likelihood fit, as discussed above.

Although no immediate work is planned on this project, future efforts will have to address the problem of formulating the DDL functional form based on the inclusion of higher order terms in the underlying maximum likelihood fit. This, unfortunately, has already proven to be a difficult task. Finally, further application to known systems will be necessary to gain additional practical exprience.

V. Summary

Research performed during this funding period has included: an improved formulation for the energy density of an infinite lattice of vortex-like objects, and an investigation of all possible geometries of the two-component, triangular vortex lattice and its energy of slip displacement; an extensive numerical investigation of the properties of an important new model of the interaction of a vortex with a bluff body in an open flow, which exhibits a 'chaotic trapping' phenomenon; identification of low-dimensional chaos in a vortex interaction experiment in open flows, utilizing a new method for determining minimum embedding dimension; an extensive new method for using time series of data from chaotic, dissipative systems to do system identification and prediction; and finally development of a new method for computing the minimum embedding dimension from a time series, used in the vortex experiment above, which is unambiguous, requires less data, and is more robust to noise than conventional techniques.

All of the above projects have resulted in significant new results regarding either vortices or vortex dominated flows; the determination or description of chaos in systems; or both aspects; and have culminated in three published papers, two papers currently in review, and one pre-print currently in preparation.

APPENDIX A

Energy of infinite vortex lattices

L. J. Campbell
Theoretical Division, MS-B262, Los Alamos National Laboratory, Los Alamos, New Mexico 87545

M. M. Doria*
Center for Nonlinear Studies, MS-B258, Los Alamos National Laboratory, Los Alamos, New Mexico 87545

J. B. Kadtke
Institute for Nonlinear Science, University of California San Diego, La Jolla, California 92093
(Received 28 November 1988)

An expression is derived for the energy density of a lattice of point vortices (or other logarithmic objects) having an arbitrary number of vortices of arbitrary strengths in an arbitrary unit cell. The result is expressed in the form of a rapidly convergent series well suited for numerical evaluation. The effects of separately changing the shape and dimensions of the unit cell are shown for simple cases, and the energy of the triangular lattice is calculated as a function of slip displacement.

I. INTRODUCTION

We consider the problem of finding the energy of an infinite number of classical point particles confined to a planar lattice and interacting pair wise with a logarithmic potential. These particles will be viewed as vortices in an Eulerian fluid; they are also equivalent to rectilinear line charges, line currents, or screw dislocations. Our objective is to find the relative energy of different configurations of J vortices having strengths $\Gamma_1,\Gamma_2,\ldots,\Gamma_J$ in a unit cell defined by the lengths L_1 and L_2 of its sides and the angle ϕ between them.

If the sum of the vorticity strengths is not zero in the unit cell the system is stationary only in a coordinate frame rotating with angular velocity Ω,

$$\Omega=\frac{\Gamma}{2L_1L_2\sin\phi},\quad \Gamma\equiv\sum_{j=1}^{J}\Gamma_j\ .$$

We consider the lattice only in such a frame or, equivalently, in a nonrotating frame with an imposed background solid-body rotation of the opposite sign. $-\Omega r$. Similarly, an opposite uniform background charge or current would be needed for line charges or currents. Such constant background fields play no role in the lattice properties, and serve merely to cancel formal[1] singularities that occur at zero wave number. Of course, these background fields must be explicitly included to study the global properties of finite[2] systems.

The task of deriving lattice sums for Coulomb interactions has a long history.[3] Our purpose here is to obtain the most efficient lattice sum for a general two-dimensional lattice and our method based on results by Glasser,[4] who considered the particular case of a rectangular unit cell ($\phi=90°$). In addition to obtaining a rapidly convergent lattice summation, we obtain an expression for the energy density of a vortex lattice that is invariant to physically equivalent designations of the unit cell, which are not necessarily primitive cells. By means of this expression it becomes easy to compare the energy of *all possible* lattices containing the same mixture of vortex species.

II. LATTICE ENERGY

The total energy due to mutual vortex interaction is

$$E_T=-\frac{d}{4\pi}\sum_i{\sum_j}'\Gamma_i\Gamma_j\ln|\mathbf{r}_i-\mathbf{r}_j|\ , \tag{1}$$

where d is the fluid density (mass per unit area) and the double sum omits $i=j$. For an infinite lattice E_T is unbounded, even in the presence of a background. However this unboundedness is easily avoided by considering the energy per vortex E, which is finite:

$$E=\lim_{M\to\infty}\frac{4\pi}{dJM}E_T\ , \tag{2}$$

where M is the number of unit cells. It is convenient to subdivide the sum over all vortices into sums over the J vortex species in all unit cells,

$$\sum_j=\sum_{j_1}\sum_{j_2}+\cdots+\sum_{j_J} \tag{3}$$

and to note that

$$\sum_{j_a}{\sum_{j_b}}'\Gamma_a\Gamma_\beta\ln|\mathbf{r}_{j_a}-\mathbf{r}_{j_\beta}|=M\Gamma_a\Gamma_\beta{\sum_n}'\ln|\mathbf{r}_\alpha^0-\mathbf{r}_\beta^0+\mathbf{L}_n|\ . \tag{4}$$

The sum in the above equation is over all integers $n_1,n_2=0,\pm1,\pm2,\ldots$, except if $\alpha=\beta$, in which case $n_1=n_2=0$ must be omitted. The vortex positions are $\mathbf{r}_\alpha^0$, $\alpha=1,\ldots,J$ in a reference unit cell and

$$\mathbf{L}_n=L_1n_1\mathbf{e}_1+L_2n_2\mathbf{e}_2 \tag{5}$$

is a generic lattice vector ($\mathbf{e}_1\cdot\mathbf{e}_2=\cos\phi$). Using Eqs. (3) and (4) in Eq. (2) gives

 ©1989 The American Physical Society

$$E=-\overline{\Gamma}^2\sum_n{}'\ln|\mathbf{L}_n|-\frac{2}{J}\sum_{\substack{\alpha=1\\ \alpha<\beta}}^{J-1}\sum_{\beta=2}^{J}\Gamma_\alpha\Gamma_\beta\sum_n\ln|\mathbf{r}^0_{\alpha\beta}+\mathbf{L}_n|\ , \tag{6}$$

where

$$\overline{\Gamma}^2\equiv\frac{1}{J}\sum_{\alpha=1}^{J}\Gamma_\alpha^2 \tag{7}$$

and

$$\mathbf{r}^0_{\alpha\beta}=\mathbf{r}^0_\alpha-\mathbf{r}^0_\beta\ .$$

III. LATTICE SUM

To evaluate the lattice sums in Eq. (6) express the Fourier transform of the logarithm function using box normalization,

$$-\ln|\mathbf{x}|=\lim_{\mu\to0}\frac{2\pi}{s_1s_2}\sum_{\mathbf{k}}\frac{\exp(i\mathbf{k}\cdot\mathbf{x})}{k^2+\mu^2},\quad \mathbf{k}=2\pi\left[\frac{n_1}{s_1},\frac{n_2}{s_2}\right], \tag{8}$$

in the limit where s_1 and s_2 become infinite. A nonzero "mass" parameter μ changes the logarithm function into a short-ranged one, and is a fundamental parameter for understanding the effect of the background. To perform the lattice sum it is convenient to employ the so-called reciprocal-lattice vectors g defined by

$$\mathbf{g}=\frac{2\pi}{\sin\phi}\left[\frac{m_1}{L_1}\mathbf{v}_1+\frac{m_2}{L_2}\mathbf{v}_2\right],\quad \mathbf{v}_i\cdot\mathbf{e}_j=\delta_{ij}\sin\phi\ . \tag{9}$$

Then,

$$-\sum_n\ln|\mathbf{x}+\mathbf{L}_n|=V(x)+c_\infty\ , \tag{10}$$

where

$$V(x)=\frac{2\pi}{L_1L_2\sin\phi}\sum_{\mathbf{g}(\neq0)}\frac{e^{i\mathbf{g}\cdot\mathbf{x}}}{g^2}\ , \tag{11}$$

$$c_\infty=\lim_{\mu\to0}\frac{2\pi}{\mu^2L_1L_2\sin\phi}\ . \tag{12}$$

The divergent constant c_∞ corresponds to the $\mathbf{g}=0$ component ($m_1=m_2=0$); the effect of the background is to cancel this divergent constant.

To apply Glasser's method one first writes Eq. (11) in a more explicit form,

$$V(x)=\frac{\rho\sin\phi}{2\pi}\sum_{m_1,m_2(\neq0)}\frac{e^{2\pi i(m_1y_1+m_2y_2)}}{m_1^2+m_2^2\rho^2-2m_1m_2\rho\cos\phi}\ , \tag{13}$$

where $y_1=(x_1\sin\phi-x_2\cos\phi)/L_1\sin\phi$ and $y_2=x_2/L_2\sin\phi$. The same sequence of transformations of Ref. 4 then leads to

$$V(x)=(\sin\phi/\rho\pi)\sum_{k=1}^{\infty}\cos(kz_2/\sin\phi)/k^2-\tfrac{1}{2}\ln\prod_{s=-\infty}^{\infty}h(s,z_1,z_2)\ , \tag{14}$$

where

$$h(s,z_1,z_2)=1-2e^{-|z_2+2\pi s\sin\phi|/\rho}\cos\left[z_1+\frac{2\pi s}{\rho}\cos\phi\right]+e^{-2|z_2+2\pi s\sin\phi|/\rho}\ , \tag{15}$$

with $z_i\equiv2\pi x_i/L_i$ and $\rho\equiv L_1/L_2$.

In terms of these new variables a lattice translation $\mathbf{x}\to\mathbf{x}+\mathbf{L}_n$ becomes $z_1\to z_1+2\pi n_1+2\pi n_2\cos\phi/\rho$ and $z_2\to z_2+2\pi n_2\sin\phi$. It is easy to verify that Eq. (14) is invariant under lattice translations, consistent with Eq. (10). The first summation of Eq. (14) can be performed, giving the more efficient representation,[5]

$$(\sin\phi/\pi)\sum_{k=1}^{\infty}\cos(kz_2/\sin\phi)/k^2=|z_2|(|z_2|/\sin\phi-2\pi)/4\pi+\pi(\sin\phi)/6\ , \tag{16}$$

valid for $|z_2|\le2\pi\sin\phi$. The consequent loss of translational invariance in the $\mathbf{e}_2$ direction causes no difficulty in numerical evaluations.

The expression for $V(x)$ given by Eqs. (14)–(16) converges quite rapidly. In practice, the evaluation of the infinite product of terms $h(s,z_1,z_2)$ reduces to the multiplication of about four to eight terms because, for large integers s, $h(s,z_1,z_2)$ is dominated by unity plus terms proportional to $\exp(-s)$, which have a very fast decay. This product expansion is almost identical to the expansion for the Jacobi Θ functions; for the special case of $\phi=90°$ studied by Glasser, it reduces to them. Finally, an expression is needed for the first term of Eq. (6), which is the energy of identical vortices on a primitive lattice, a result derived also by Tkachenko.[6] This term is equivalent to the following limit:

$$\sum_n{}'\ln|\mathbf{L}_n|=\lim_{x\to0}\left\{\sum_n\ln|\mathbf{x}+\mathbf{L}_n|-\ln|\mathbf{x}|\right\}\ . \tag{17}$$

Performing this limit on Eq. (14) gives

$$-\sum_n{}'\ln|\mathbf{L}_n|=\frac{\pi}{6\rho}\sin\phi-\ln(2\pi/L_1)-\ln\prod_{s=1}^{\infty}\{1-2e^{-2\pi s|\sin\phi|/\rho}\cos[2\pi s(\cos\phi)/\rho]+e^{-4\pi s|\sin\phi|/\rho}\}\ . \tag{18}$$

Now we scale the energy to obtain equal energies for physically equivalent lattices. This is simply done by noting that E in Eq. (6) is the energy per vortex. Hence, scaling the lengths L_1 and L_2 by a constant α gives the correct normalized energy and renders a constant vortex density. This causes no changes to the ratio L_1/L_2, but in Eq. (18) the di-

mensional constant L_1 enters alone. We choose the density to be unity, i.e.,

$$\frac{J}{\alpha^2 L_1 L_2 \sin\phi} \equiv 1 \ . \tag{19}$$

Solving for α and multiplying L_1 in Eq. (18) by α then gives

$$\ln(2\pi/\alpha L_1) = \ln\left[2\pi\left[\frac{\sin\phi}{J\rho}\right]^{1/2}\right], \tag{20}$$

which removes all dimensional constants from Eq. (6), except for the Γ_α, whose dimensions are trivial to remove.

The final result for the energy density is

$$E = \overline{\Gamma}^2\left\{\frac{1}{\rho}\frac{\pi}{6}\sin\phi - \ln\left[2\pi\left[\frac{\sin\phi}{J\rho}\right]^{1/2}\right] - \ln\prod_{s=1}^{\infty} h(s,0,0)\right\}$$
$$+\frac{2}{J}\sum_{\substack{i=1\\ i<j}}^{J}\Gamma_i\Gamma_j\left\{\frac{1}{\rho}\left[\frac{|z_{2,ij}|}{4\pi}\left[\frac{|z_{2,ij}|}{\sin\phi}-2\pi\right]+\frac{\pi}{6}\sin\phi\right]-\tfrac{1}{2}\ln\prod_{s=-\infty}^{\infty} h(s,z_{1,ij},z_{2,ij})\right\}, \tag{21}$$

where $z_{1,ij}=2\pi(\mathbf{r}_i^0-\mathbf{r}_j^0)\cdot\hat{x}/L_1$, $z_{2,ij}=2\pi(\mathbf{r}_i^0-\mathbf{r}_j^0)\cdot\hat{y}/L_2$, and $h(s,z_1,z_2)$ is defined in Eq. (15). This expression for E gives the relative energy density of lattices containing fixed ratios of vortex species having fixed strengths. To compare the energies of lattices which do not have the same mixtures of vortices requires assumptions or physical information about the vortex self-energies.

What makes Eq. (21) useful for numerical evaluation is the fast convergence of the function $h(s,z_1,z_2)$. Some applications, not discussed here, require calculating the partial derivatives of E, for which it is convenient to change the unit-cell variables ρ and ϕ to $\sigma=2\pi(\sin\phi)/\rho$ and $\chi=2\pi(\cos\phi)/\rho$.

IV. EXAMPLES

The foregoing results will now be applied to some simple examples. First, consider the change of lattice energy density induced by varying the angle ϕ between the lattice generators while holding fixed the lengths of the unit cell and the relative positions of the vortices. Three cases will be considered: (a) one vortex per unit cell with $L_1=L_2$; (b) two vortices per unit cell, also with $L_1=L_2$; and finally (c) two vortices per unit cell with $L_1=L_2/\sqrt{3}$. The results as calculated from Eq. (21) are shown in Fig. 1. The triangular lattice occurs for (a) when $\phi=60°$ and 120° and for (c) when $\phi=90°$. The

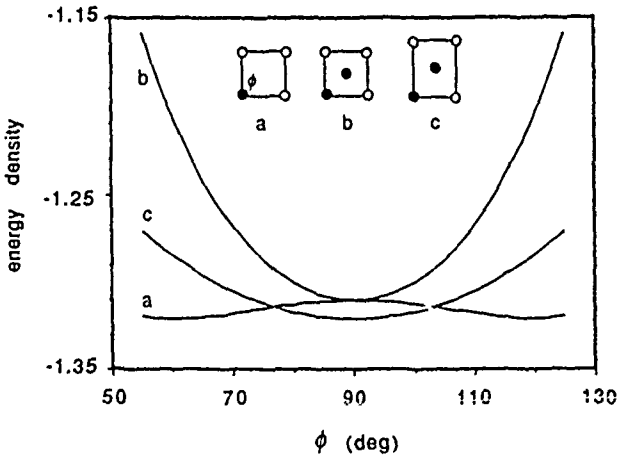

FIG. 1. Effect of varying the angle ϕ between the unit-cell generators for fixed unit-cell lengths. The different unit cells are illustrated for $\phi=90°$. (a) One vortex per unit cell with $L_1=L_2$. (b) Two vortices at positions (0,0) and (0.5,0.5) with respect to the unit-cell lengths, $L_1=L_2$. (c) Two vortices at positions (0,0) and $(0,\sqrt{3}/2)$ in a unit cell with $L_1=L_2/\sqrt{3}=1$. The energy density is the energy per vortex in units of $d\Gamma^2/4\pi$, where d is the fluid density and Γ is the unit of circulation.

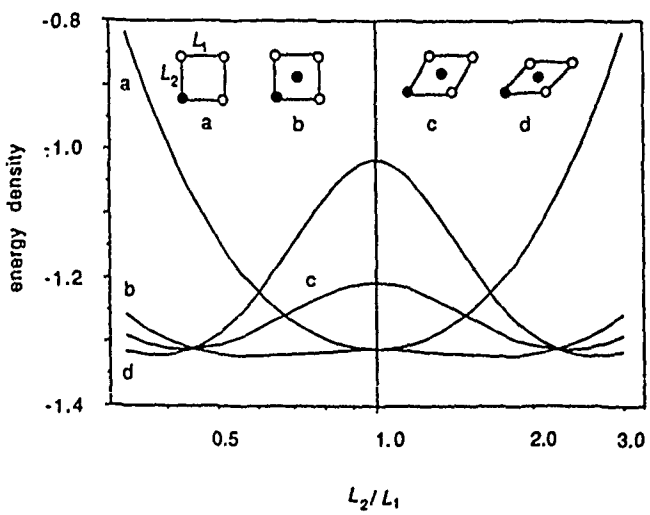

FIG. 2. Effect of changing the aspect ratio L_2/L_1 for fixed angle ϕ. The various unit cells are illustrated for $L_2/L_1=1$, with the vortices associated with the unit cell indicated by solid circles. (a) One vortex with $\phi=90°$. (b) Two vortices with $\phi=90°$. (c) Two vortices with $\phi=60°$. (d) Two vortices with $\phi=45°$.

square lattice occurs for both (a) and (b) at $\phi=90°$. The energy densities of the triangular and square lattices are $-1.321\,117\,428\,4$ and $-1.310\,532\,925\,9$, respectively. (Earlier evaluations of the energies of these simple lattices are equivalent within constants.[6,2]) Although curve (b) has a minimum at $\phi=90°$ this is a constrained minimum and does not result in a stable lattice; indeed, it joins curve (a) which leads to the absolute minimum.

Next, the angle ϕ is constrained and the ratio of unit-cell lengths L_2/L_1 is varied. These results are shown in Fig. 2, where the various cases are (a) one vortex and (b) two vortices per unit cell with $\phi=90°$, (c) two vortices with $\phi=60°$, and (d) two vortices with $\phi=45°$. Only curve (b) achieves the triangular lattice. This occurs at $L_2/L_1=\sqrt{3}$ and $1/\sqrt{3}$. Note that the horizontal scale is logarithmic, to illustrate the symmetry around $L_2/L_1=1$. It appears that curve (d) may also reach the low energy of the triangular lattice. In fact, it does not, nor is the minimum it does reach an unconstrained minimum of the lattice. Also despite appearances, curves (b), (c), and (d) do not mutually intersect.

Finally, the slip strength of the triangular vortex lattice is calculated for displacements along one of the principal axis directions. That is, the energy density is evaluated as a function of a rigid displacement, through one lattice spacing, of a number n of lattice rows with respect to the same number of fixed rows. The pattern repeats, of course, to infinity. During this displacement the unit-cell dimensions and angle are held fixed so, in particular, there is no change in volume. Figure 3 shows the results for various n, as labeled. Obviously, the maximum occurs for a displacement halfway between equilibrium positions and is largest for alternating single rows ($n=1$). This energy is just that of a rectangular lattice with $L_2/L_1=(\sqrt{3}/2)^{\pm1}=(0.866)^{\pm1}$, which can be verified by comparing the maximum in Fig. 3 with curve (a) in Fig. 2 at that ratio. The curves are approximately related to each other by

$$n_j[E_j(d)-E_t]\approx n_k[E_k(d)-E_t]\,, \tag{22}$$

where E_t is the triangular lattice energy density (given above) and d is the displacement. Future publications will treat other applications, especially those that involve seeking minima of the energy density in the presence of additional dynamics, mixtures of vortex strengths, and the unconstrained space of lattice variables.[7]

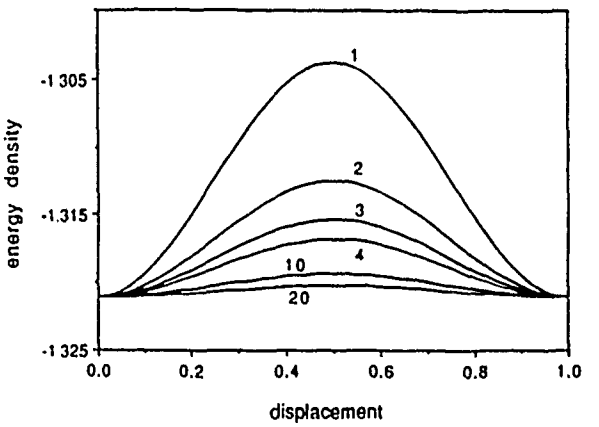

FIG. 3. Slip strength of triangular vortex lattice for the rigid displacement of n rows of vortices with respect to n stationary rows in each unit cell. The curves are labeled by n and the range of displacement is one lattice spacing along a principal axis.

V. CONCLUSION

Lattices of nonneutral vortices, like charges, have a long-range interaction which leads to a formal singularity when the lattice energy is calculated as the $N\to\infty$ limit of a finite system. By the same method as used for charges, this singularity can be removed by adding a neutralizing background. For vortices, this background is taken to be uniform, with the result that there is no phenomenon of screening. Also, like charges, the field for each vortex leads to a formal singularity in the self-energy in the limit of vanishing core size. This singularity, too, is irrelevant, except that it prevents, in the absence of further assumptions or physical information, a comparison of the lattice energies of vortex systems containing different mixtures of vortex strengths.

The energy density of the general vortex lattice (arbitrary unit cell and arbitrary number, magnitudes and signs of strengths of vortices per unit cell) is given by Eq. (21), which has the virtue of being easily evaluated numerically, in the sense of rapid convergence of its infinite products. This expression provides a new, practical tool for studying a wide range of vortex lattice problems.

ACKNOWLEDGMENTS

One of the authors (J.B.K.) wishes to acknowledge the support from the Air Force of Scientific Research Grant No. AFOSR-89-0072.

*Present address: Departamento de Física Matemática, Instituto de Física, Universidade de São Paulo, Caixo Postal 20516, Codigo de Enderecamento Postal 01498, São Paulo, São Paulo, Brazil.

[1]A. L. Fetter and J. D. Walecka, *Quantum Theory of Many-Particle Systems* (McGraw-Hill, New York, 1971), Secs. 12 and 30.

[2]L. J. Campbell and R. M. Ziff, Phys. Rev. B 20, 1886 (1979).

[3]P. P. Ewald, Ann. Phys. (N.Y.) 64, 253 (1921).

[4]M. L. Glasser, J. Math. Phys. 15, 188 (1974).

[5]I. S. Gradshtein and I. M. Ryzhik, *Table of Integrals, Series, and Products*, 4th Ed. (Academic, 1980), Sec. 1.443, p. 39.

[6]V. K. Tkachenko, Zh. Eksp. Teor. Fiz. 49, 1875 (1965) [Sov. Phys.—JETP 22, 1282 (1966)].

[7]L. J. Campbell, in *Mathematical Aspects of Vortex Dynamics*, edited by Russel E. Caflisch (Society for Industrial and Applied Mathematics, Leesburg, VA, 1989), pp. 195–204.

APPENDIX B

Chaotic vortex-body interaction

E. A. Novikov
Institute for Nonlinear Science, R-002
University of California, San Diego
La Jolla, California 92093

ABSTRACT

It is shown, by using the Poincaré-Melnikov-Arnold method, that the motion of a linear vortex in the flow past a cylindrical body is chaotic. In particular, a vortex can be captured by the body and then, after some complicated rotation near the body, can be lost. More general problems of vortex-body interaction are discussed qualitatively. Possible applications of the theory are indicated.

The study of the chaotic interactions of linear vortices has a twelve year history with the participation of many authors (see Ref. 1-15 and references therein). Generalization to the axisymmetric flows have been indicated in Ref. 16. The main conclusion from these studies is that two-dimensional and axisymmetrical flows of ideal incompressible fluid are generally non-integrable and with appropriate initial conditions exhibit chaos. This is in contrast with the opinion which prevailed earlier among the advocates of the soliton approach to the hydrodynamics. The chaotic motion of vortices has various applications, in particular, to the problem of weather prediction..[1,17]

The main goal of this note is to indicate novel features of chaotic motion which arise in the presence of a moving body. Firstly, it is enough to have only one vortex in order to get chaotic motion. Secondly, the mechanism of generation of chaos is very transparent in the vortex-body system. Thirdly, we get new phenomena — chaotic capture — loss of vortex by a moving body.

Vortex-body interactions are important in many situations. The most dramatic examples are the aviacatastrophies, caused by a vortex initiated by downdraft of cold air,[18] and the destruction of buildings by tornado. We will start with an analytical description of the motion of a linear vortex in the flow past a circular cylinder. Then we will make some qualitative remarks about more general problems.

In the frame of reference moving with the cylinder, the velocity of a linear vortex in an ideal fluid is a Hamiltonian superposition of two parts of motion. The first part corresponds to the potential motion of fluid relative to the cylinder (see, for example, Ref. 19), such as if the vortex has zero intensity. The second part of the motion is induced by the interaction of vortex with cylinder. In the case of circular cylinder, this motion is induced by an *image* vortex, placed inside the cylinder at the distance from the center $\bar{r} = a^2/r$, where a is the radius of cylinder and r corresponds to the position of vortex. Both parts of motion have zero normal components of velocity at the surface of the cylinder.

We will scale distances by a and time by a/u_0, where u_0 is the characteristic fluid velocity at infinity. If the fluid velocity at infinity is constant, then the problem is characterized by only one nondimensional parameter $\sigma = \kappa/2\pi a u_0$, where κ is the vortex intensity. In the case of vibration of cylinder, we have the relative fluid velocity at infinity

$$u(t) = u_0(1 + \epsilon \sin \omega t), \tag{1}$$

where ϵ and ω are the nondimensional amplitude and frequency of vibration.

In polar coordinates (ρ, ϕ) with origin in the center of the cylinder, the Hamiltonian system for the motion of the vortex has the form

$$\frac{d\rho}{dt} = \frac{1}{\rho}\frac{\partial H}{\partial \phi} = -\sin\phi\left(1 - \frac{1}{\rho^2}\right)(1 + \epsilon \sin \omega t), \tag{2}$$

$$\frac{d\phi}{dt} = -\frac{1}{\rho}\frac{\partial H}{\partial \rho} = -\frac{\cos\phi}{\rho}\left(1 + \frac{1}{\rho^2}\right)(1 + \epsilon \sin \omega t) - \frac{\sigma}{\rho^2 - 1}, \tag{3}$$

$$H = H_0 + \epsilon H_1, \qquad \rho = \frac{r}{a} > 1$$

$$H_0 = \frac{\cos\phi}{\rho}(\rho^2 - 1) + \frac{\sigma}{2}\ln(\rho^2 - 1), \tag{4}$$

$$H_1 = \frac{\cos\phi}{\rho}(\rho^2 - 1)\sin\omega t. \tag{5}$$

With $\epsilon = 0$, the system (2)-(4) has a general analytical solution. The vortex trajectories (4) are presented in Fig. 1. Without loss of generality, we assume that the motion of fluid around the vortex is clockwise ($\sigma < 0$) and the direction of fluid velocity at infinity is from the top to the bottom on Fig. 1. We see that there is a homoclinic vortex trajectory $[\rho_0(t), \phi_0(t)]$ with the hyperbolic stationary point at $\phi = 0, \rho = \rho_*, (\sigma = \rho_*^{-3} - \rho_*)$. The homoclinic trajectory separates the region where the vortex is captured by the body.

The stationary trajectories of vortex in the flow past circular cylinder have been studied in Ref. 20 without recognizing the Hamiltonian structure of the problem. The homoclinic trajectory have not been indicated in Ref. 20, probably because at that time the homoclinic trajectory was considered as something pathological and physically irrelevant to the problem. Now we know how important the homoclinic trajectories are for the generation of chaotic motion.

For ϵ small but non-zero, the system (2), (3) has no analytic integrals of motion. It possesses transversal intersecting stable and unstable manifolds — namely, the Poincaré maps $\mathcal{P}(t_0)$ which advance a solution by one period $T = 2\pi/\omega$ starting at time t_0, possess transversal homoclinic points. We will show this by using the Poincaré-Melnikov-Arnold (PMA) method. This type of behavior of dynamical systems is called *chaotic*.

According to PMA method,[21,22,23] we consider the *Melnikov function*

$$M(t_0) = \int_{-\infty}^{\infty} \{H_0, H_1\}(\rho_0(t - t_0), \phi_0(t - t_0); t)dt, \tag{6}$$

where $\{,\}$ denote the Poisson brackets

$$\{A, B\}(\rho, \phi; t) = \frac{1}{\rho}\left[\frac{\partial A(\rho,\phi,t)}{\partial \rho}\frac{\partial B(\rho,\phi,t)}{\partial \phi} - \frac{\partial A(\rho,\phi,t)}{\partial \phi}\frac{\partial B(\rho,\phi,t)}{\partial \rho}\right].$$

Integral (6) is taken along the homoclinic trajectory $[\rho_0(t - t_0), \phi(t - t_0)]$. We will show that $M(t_0)$ has simple zeros.

From (4)-(6), by a change of variable, we get:

$$\begin{aligned} M(t_0) &= -\sigma \int_{-\infty}^{\infty} \frac{\sin[\phi_0(t - t_0)]}{\rho_0(t - t_0)} \sin\omega t\, dt \\ &= -\sigma \int_{-\infty}^{\infty} \frac{\sin[\phi_0(t)]}{\rho_0(t)}(\sin\omega t \cos\omega t_0 + \cos\omega t \sin\omega t_0)\, dt \end{aligned} \tag{7}$$

Since $\phi_0(t) \to 0, \rho_0(t) \to \rho_*$ exponentially near the hyperbolic stationary point (when $t \to \pm\infty$), the integral (7) is convergent. It is convenient to choose the initial position: $\phi_0(0) = \pi$. In this case we have $\phi_0(-t) = -\phi_0(t), \rho_0(-t) = \rho_0(t)$ and (7) reduces to:

$$M(t_0) = -\sigma \cos\omega t_0 \int_{-\infty}^{\infty} \frac{\sin[\phi_0(t)]}{\rho_0(t)} \sin\omega t\, dt. \tag{8}$$

The integral in (8) is not identically zero, because it is the Fourier transform of a function which is not identically zero. Function $M(t_0)$ clearly has simple zeros. According to PMA theory,

this proves that system (2,3) has no analytic integrals of motion and vortex trajectory is chaotic. In particular, the vortex which is initially far from the body, can intersect the homoclinic loop and will be captured by the body. After several complicated revolutions around the body, the vortex will eventually escape.[24] In connection with the capture-loss phenomena, we have the following theorem. Let S be the set of positions of the vortex at time t_0, for which the vortex for all $t \geq t_0$ will stay inside a circle C surrounding the body. It can be proven that the subset of S, for which the vortex was outside of C of some $t < t_0$, has zero measure. The proof is the same as in the Littlewood's theorem[25] for conservative (gravitational) system. The only condition which matters is the preservation of volume in phase space. The nonautonomous Hamiltonian system clearly satisfies this condition. The theorem remains true if instead of circle C surrounding the body, we choose any area in phase space.

In the case of arbitrary shape of a cylindrical body we still have Hamiltonian superposition of external and induced motion of the vortex in terms of corresponding Green's functions for the Laplace operator. The external velocity is finite everywhere. The induced velocity is infinite near the body and zero at infinity. Having this in mind, we generally can expect existence of a homoclinic separatrice with hyperbolic stationary point, where two parts of velocity are balanced (in the case of stationary external velocity). Thus, the described above phenomena of chaotic vortex-body interaction seems to be generic.

The local kinematic pressure, exerted on the surface of the body by the vortex, is of the order of κ^2/d^2, where d is the distance between vortex and body. In the chaotic regime of motion, a vortex can come closer, it will spend more time near the body and is more likely to create a destructive impact on the body.

The generalization to the case when we have additional circulation κ_0 around the cylinder is straightforward. In this case we have to add into (4) the term $-\sigma_0 \ln \rho$, where $\sigma_0 = \kappa_0/2\pi a u_0$. By using the above described procedure, it is easy to show that in the time-periodic external flow, the motion of fluid particles becomes chaotic even without a vortex ($\sigma = 0$) when $|\sigma_0| > 2$.

Condition $|\sigma_0| > 2$ insures the existence of a hyperbolic stagnation point when $\sigma = 0$. Generally, when $\sigma \neq 0$ and $\sigma_0 \neq 0$, the unperturbed system has several stagnation points (hyperbolic and elliptic). With $\epsilon \neq 0$, the vortex can chaotically change the direction of rotation during the capture. This is clear physically, but it has to be investigated in detail numerically.

In practical problems, the vortex has a finite core, which leads to a system with an infinite number of degrees of freedom. If the size of the core is small in comparison with the size of the body, the multipole representation of the vortex can be used. In the simplest representation we have two closely located concentrated vortices, which rotate around each other. In this case we get chaos even without oscillation of the cylinder ($\epsilon = 0$). The proof is lengthy and will not be presented here, but the idea is simple — the reduction of the Hamiltonian system.[23] The slow variables are the coordinates of the center of vorticity, the angle of mutual rotation of vortices plays the role of time and the distance between vortices is a small parameter. In the case of bigger distances between two vortices, one of the vortices can be captured forever and other will escape. The capture of one of the vortices does not contradict Littlewood's theorem. This kind of partial capture of vorticity in a more complex situation happens when a tornado hits a building.

The problem of three-dimensional chaotic vortex-body interactions with the effects of vortex stretching and reconnection is more difficult and profound. In this case we plan to use the method of three-dimensional solenoidal vortex singularities (vortons), which includes a mechanism of inviscid dissipation of energy.[18,26]

Acknowledgements

This work is supported by the U.S. Department of Energy under grant DE-FG030-87ER13801. I wish to thank M. Mikulska for help in preparing Figure 1.

Figure Captions

1. The vortex trajectories in the flow past a cylinder ($\rho_* = 3$).

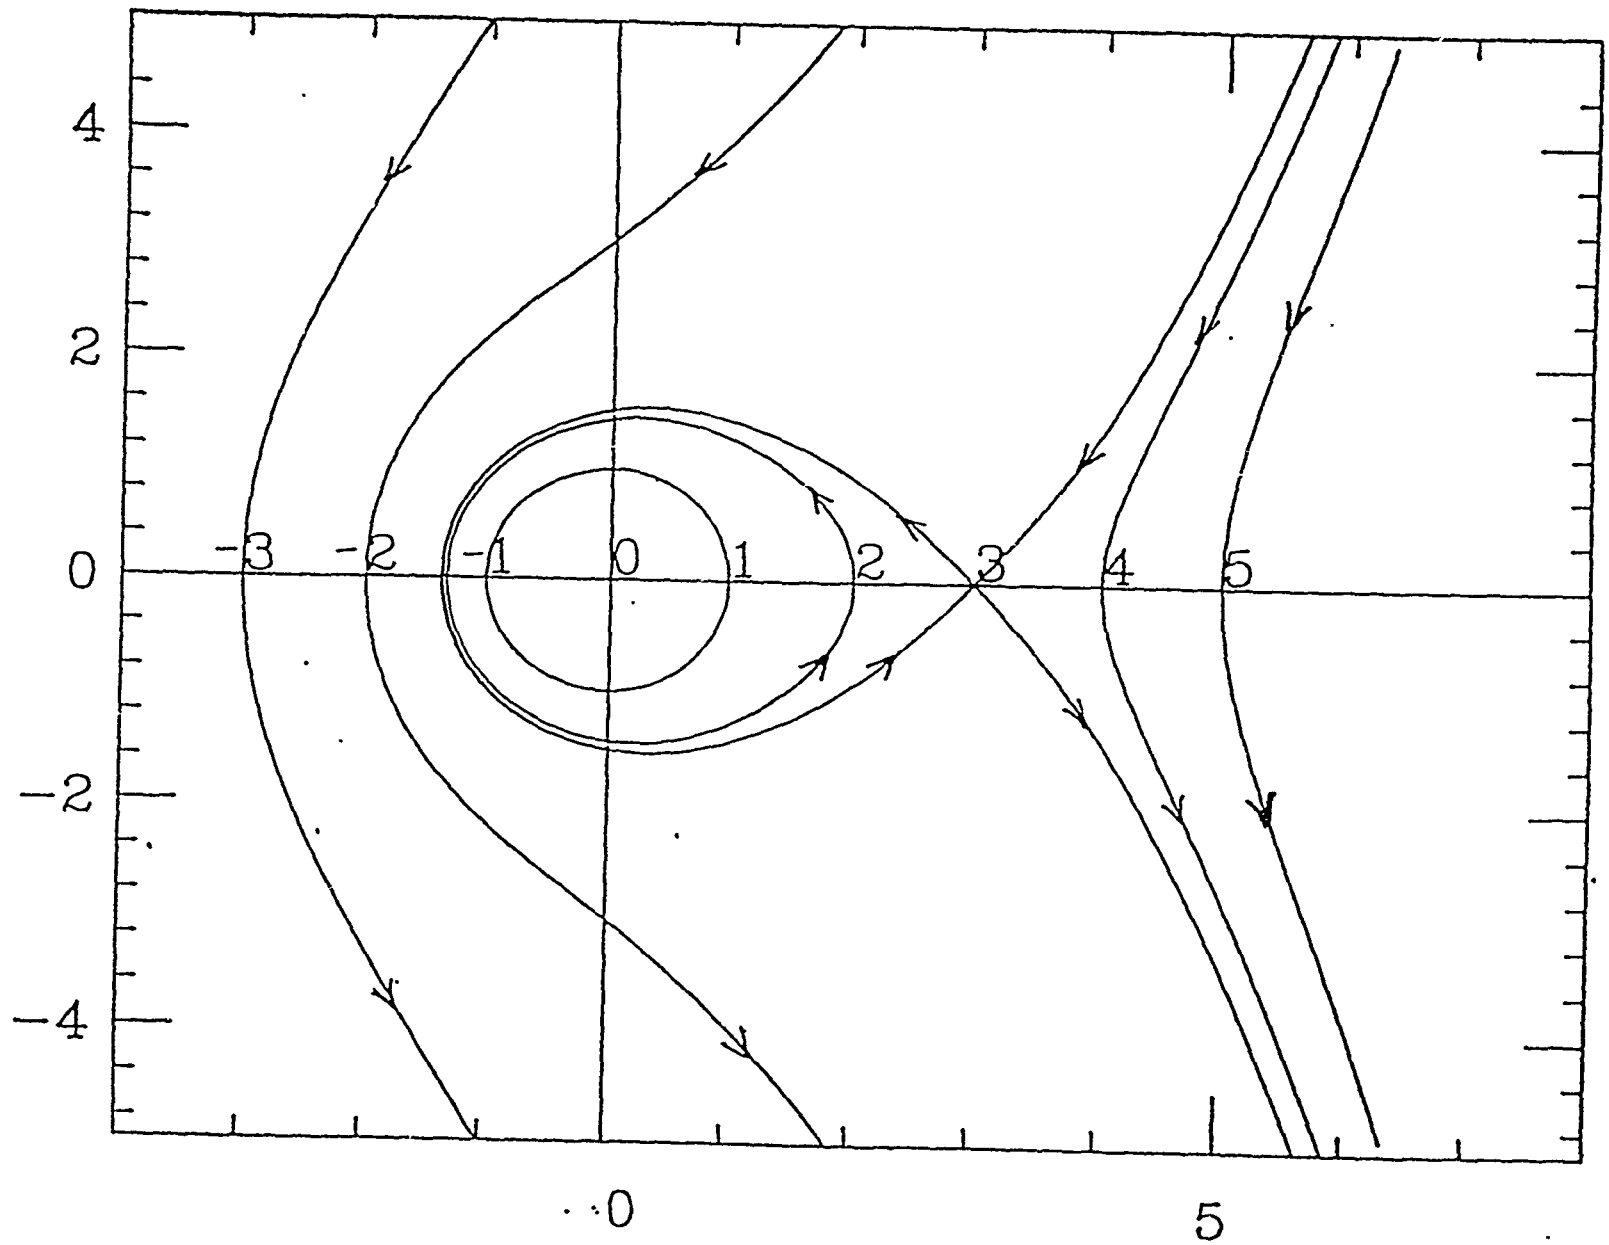
4
2
0
−2
−4
−3
−2
−1
0
1
2
3
4
5
0
5

References

[1]Novikov, E.A. and Sedov, Y.B., *Sov. Phys. JETP* 48 (3), 440 (1978).

[2]Novikov, E.A. and Sedov, Y.B., *JETP Lett.* 29 (12), 20 (1979).

[3]Inogamov, N.A. and Manakov, S.V., preprint (1979).

[4]Ziglin, S.L., *Sov. Math. Dokl.* 21, 296 (1980).

[5]Aref, H. and Pomprey, N., *Phys. Lett. A* 78 (4), 279 (1980).

[6]Novikov, E.A., *Ann. N.J. Acad. Sci.* 357, 47 (1981).

[7]Aref, H., *Ann. Rev. Fluid Mech.* 15, 345 (1983).

[8]Marsden, J. and Weinstein, A., *Physica D* 7, 305 (1983).

[9]Manakov, S. and Schur, L., *JETP Lett.* 37, 54 (1983).

[10]Hasimoto, H., Ishii, K., Kimura, Y. and Sakiyama, M., Proc. IUTAM Symp. on Turbulence and Chaotic Phenomena in Fluids, Kyoto, 1983 (North-Holland, 1984), p. 231.

[11]Koiller, J. and Carvalho, S.P., preprint (1983).

[12]Aref, H., *J. Fluid Mech.* 143, 1 (1984).

[13]Hardin, J.C. and Mason, J.P., *Phys. Fluids* 27, 1583 (1984).

[14]Kimura, Y. and Hasimoto, H., *J. Phys. Soc. Japan* 55 (1), 5 (1986).

[15]Aref, H., Kadtke, J.B., Zawadzki, I., Campbell, L.J. and Eckhardt, B., *Fluid Dyn. Res.* 3, 63 (1988).

[16]Novikov, E.A., *Phys. Fluids* 28 (9), 2921 (1985).

[17]Novikov, E.A. and Chefranov, S.G., *Phys. Atmos. Oc.* 13 (6), 414 (1977).

[18]Novikov, E.A., *Boundary Layer Meteor.* 38, 305 (1987).

[19]Landau, L.D. and Lifshitz, E.M., Fluid Mechanics (Pergamon, New York, 1959) p. 26.

[20]Walton, E.T.S., *Proc. Royal Irish Academy* 38A (3), 29 (1928).

[21]Melnikov, V.K., *Trans. Moscow Math. Soc.* 12, 1 (1963).

[22]Arnold, V., *Dokl. Akad. Nauk. SSSR* 153, 9 (1964).

[23]Guckenheimer, J. and Holmes, P., Nonlinear oscillation, dynamical systems and bifurcations of vector fields, (Springer-Verlag, 1983).

[24]This prediction has been recently confirmed in the series of numerical experiments, which were performed in collaboration with J. Kadtke and will be published elsewhere.

[25]Littlewood's miscellany (Ed. by B. Bollobas), Cambridge University Press, 1986, p. 186.

[26]Novikov, E.A., *Zh Eksp. Teor. Fiz.* 84, 975 (1983) [*Sov. Phys. JETP* 57, 566 (1983)]. Aksman, M.I., Novikov, E.A. and Orszag, S.A., *Phys. Rev. Lett.* 54, 2410 (1985). Novikov, E.A., *Phys. Lett. A* 112, 327 (1985). Aksman, M.I. and Novikov, E.A., *Fluid Dyn. Res.* 3, 239 (1988). Novikov, E.A. in *Proc. of the Sixth Symp. on Energy Engineering Sciences* (Argonne National Laboratory, Argonne, IL, 1988), p. 59.

VARIATON OF TOPOLOGY WITH σ, σ_0

o = ELLIPTIC POINT x = HYPERBOLIC POINT

CASE 1: $0 < \sigma < 9.4815$

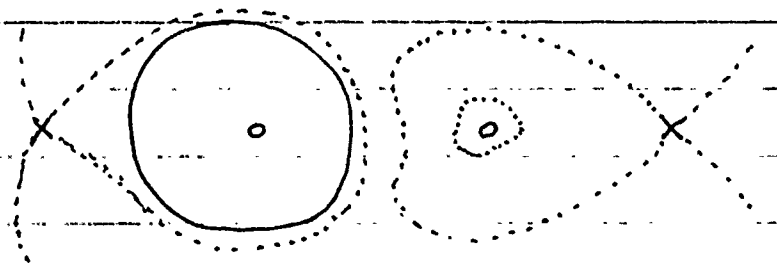

THIS JUST THE CASE WE HAVE DISCUSSED SO FAR...

CASE 2: $\sigma = 0$, $\sigma_0 = 3.3333...$

$\Rightarrow \rho_1 = 3.0$, $\rho_2 = 1.0$, $\rho_3 = -1.0$ $\rho_4 = 0.3333...$

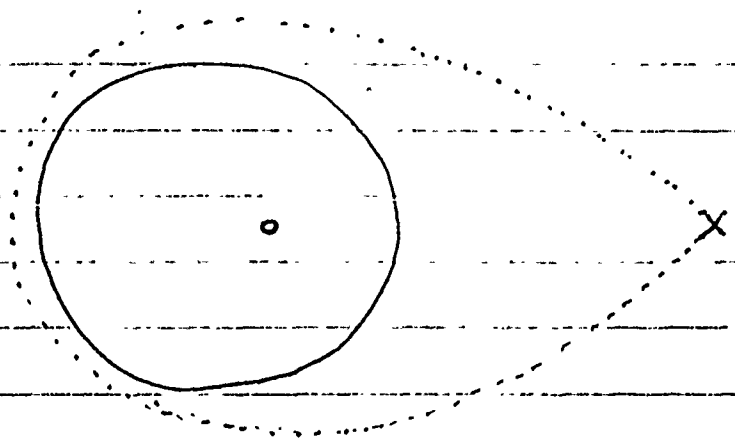

NOTE:

ROOTS AT ρ_2, ρ_3 HAVE NO APPARENT EFFECT ON FLOW FIELD.

SEE PLOT FOR ACTUAL SEPARATRIX.

CASE 2: $\sigma = 0$, $\sigma_0 = 3.333\ldots$

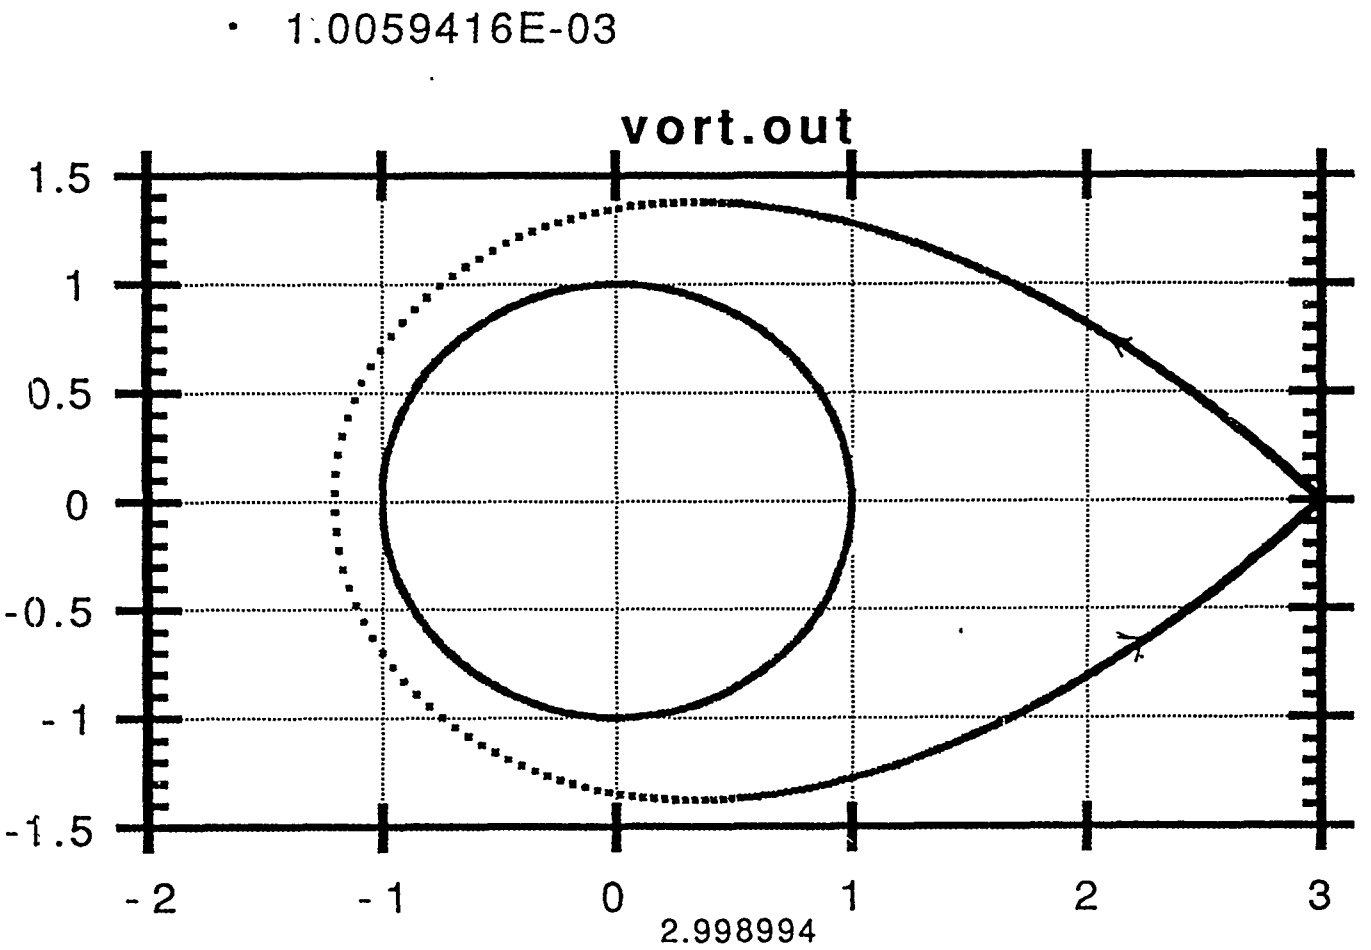

SEPARATRIX, SOLE HYPERBOLIC POINT AT (3.0, 0.0)

CASE 3: $\sigma = 9.4815$, $\sigma_0 = 14.0$ ...

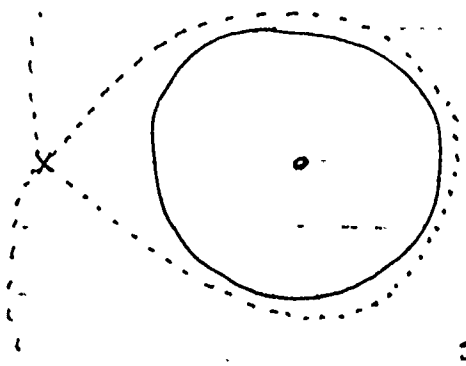

$\rho_1 = 3.0$, $\rho_2 = 3.0$

$\rho_3 = -1.55303$, $\rho_4 = 0.071545$

SEE GRAPH

POINT (DOUBLE POLE) AT 3.0 SEEMS TO HAVE ONLY ONE-HALF OF HYPERBOLIC SEPARATRIX, OTHER SIDE LOOKS ELLIPTIC!

CASE 4: $\sigma > 9.4815$

EXAMPLE PICK $\sigma = 15.0$, $\sigma_0 = 20.2083$

$\Rightarrow \rho_1 = 3.0$, $\rho_2 = 3.88957$, $\rho_3 = -1.73075$, $\rho_4 = 0.049516$

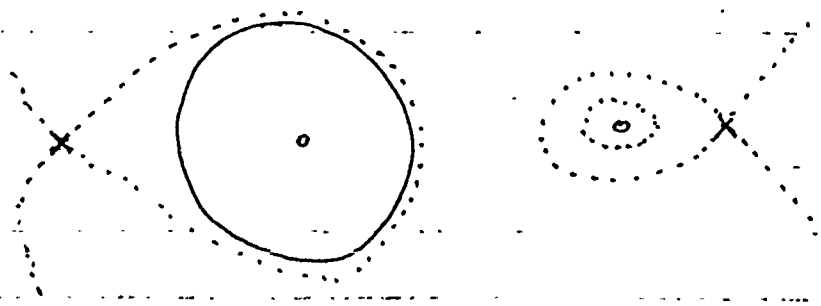

SEE PLOT

$\rho_2 \to \infty$ AS $\sigma \to \infty$

CASE 3: $\sigma = 9.4815$, $\sigma_0 = 14.0$

• 6.0296722E-04

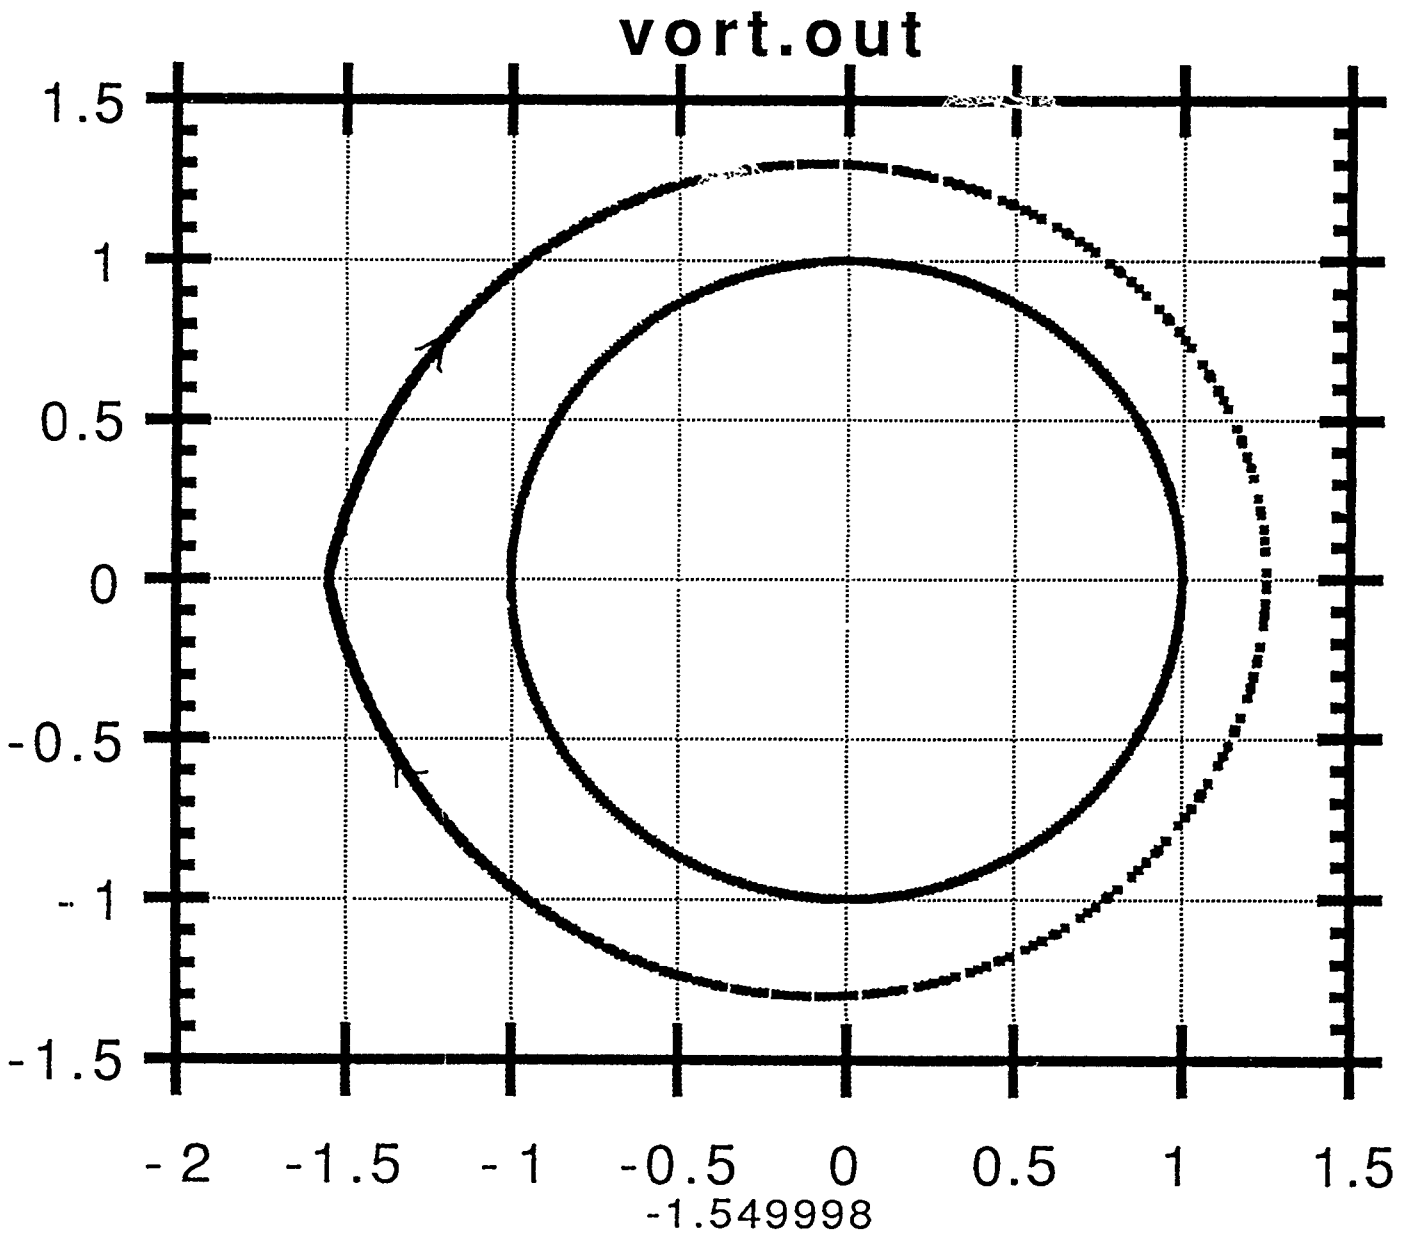

HYPERBOLIC POINT AT $\theta = -1.55303$

CASE 3: $\sigma = 9.4815$, $\sigma_0 = 14.0$

• 4.9989076E-02

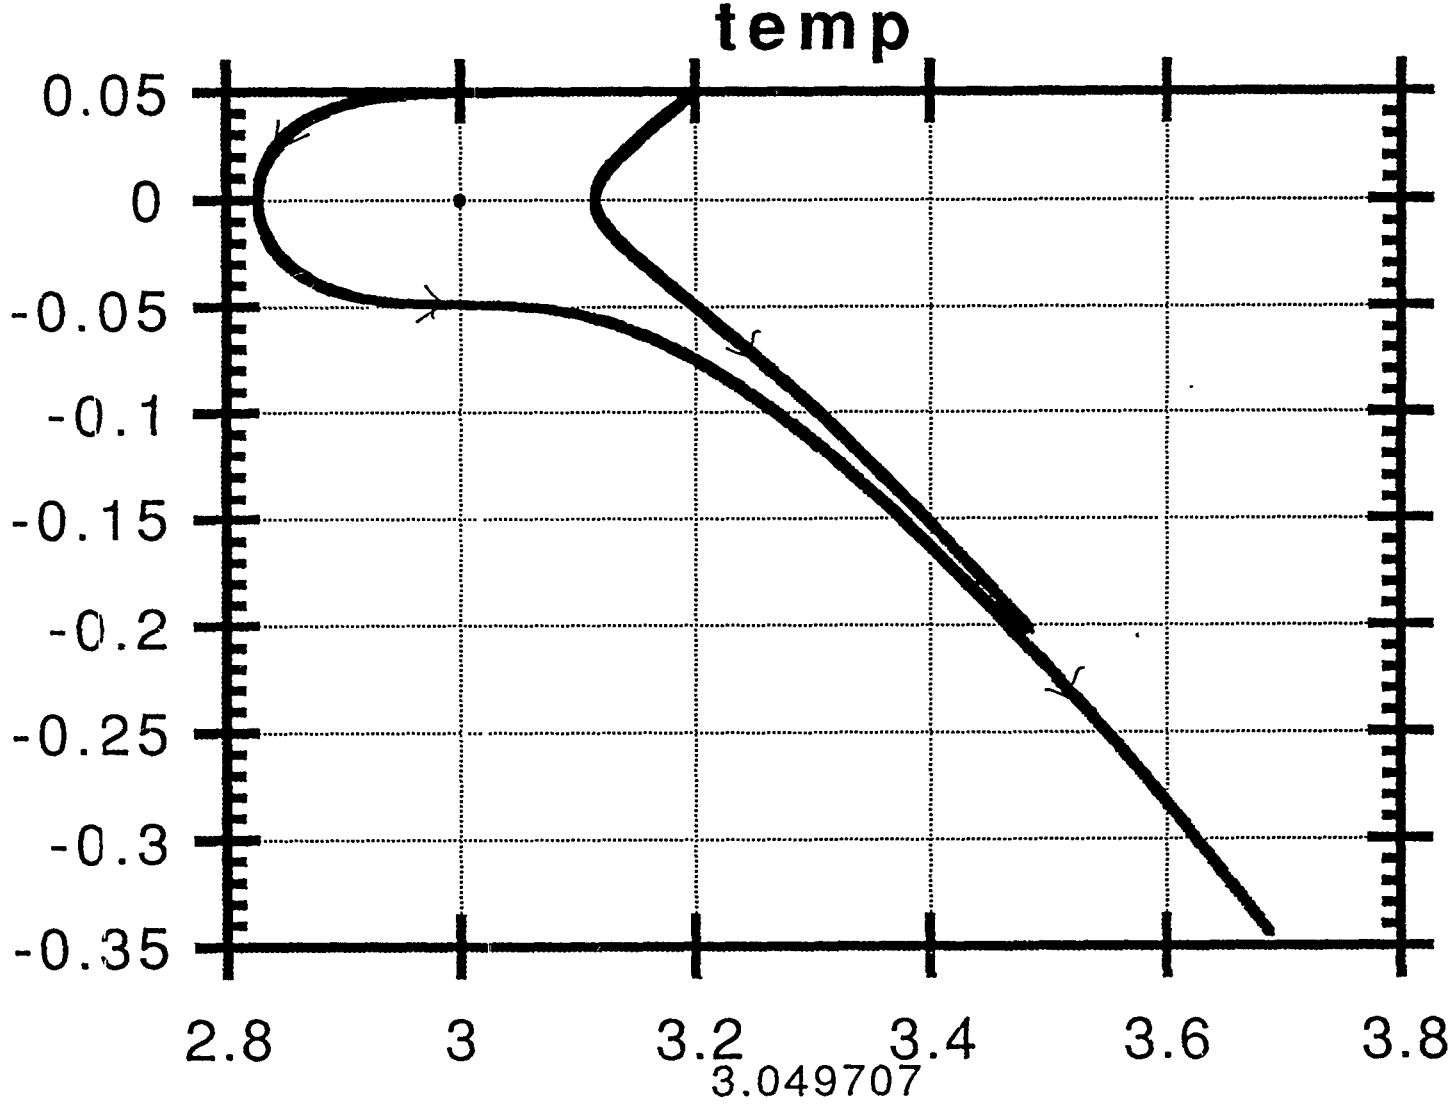

TWO ROOTS (ONE HYPERBOL., ONE ELLIP.) AT $\rho = 3.0$

CASE 4 : $\sigma = 15.0$, $\sigma_o = 20.2083$

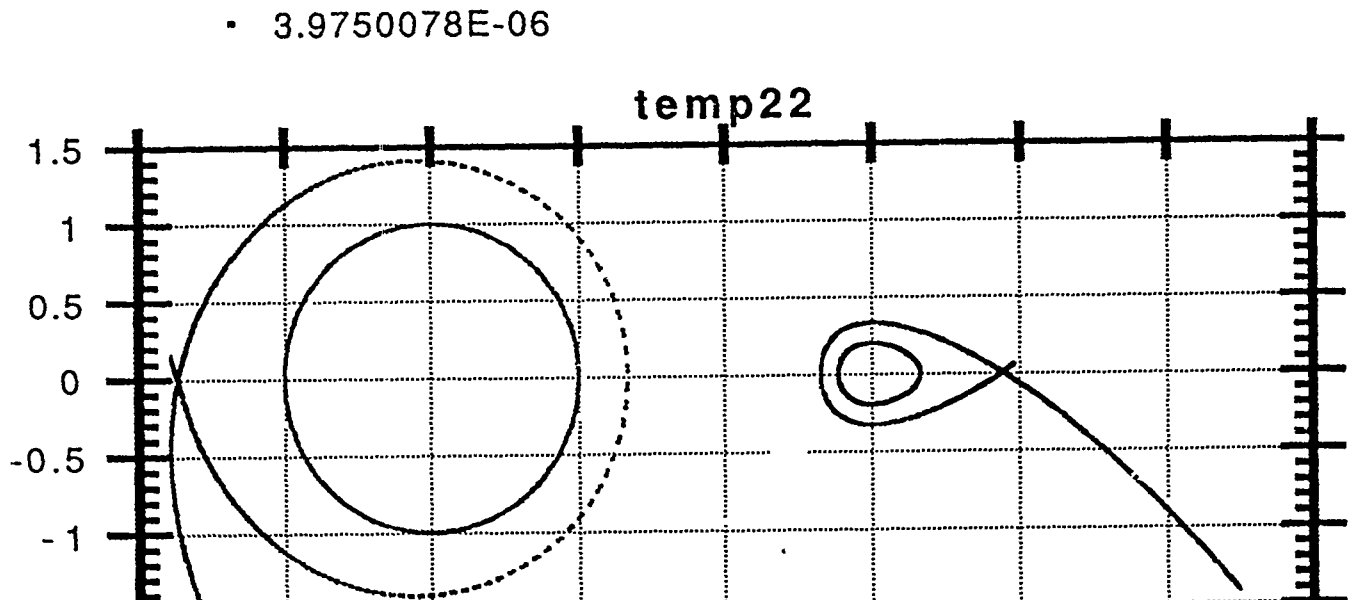

CASE 5: $\sigma = +\infty$, $\sigma_0 = +\infty$

$\Rightarrow \rho_1 = 3.0$, $\rho_2 = +\infty$, $\rho_3 = -3.0$, $\rho_4 = 0.0$

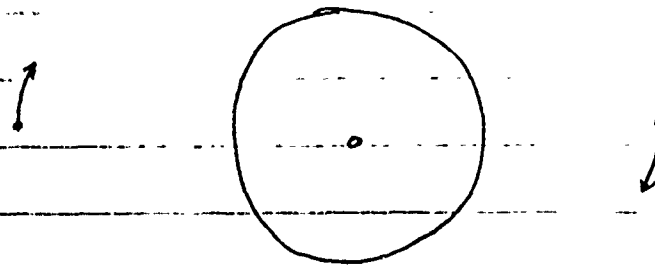

SEE PLOT

ρ_1, ρ_3 ARE NOT STAGNATION POINTS. FLUID ROTATES UNIFORMLY AROUND BODY

CASE 6: $0 > \sigma > -0.6321838$

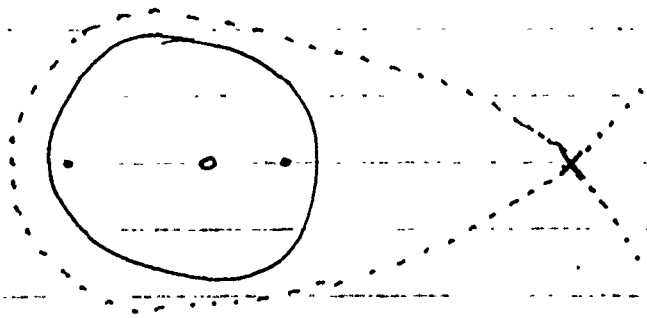

SEE PLOT

EXAMPLE:

$\sigma = -0.5$, $\sigma_0 = 2.77083$

$\rho_1 = 3.0$, $\rho_2 = 0.763154$, $\rho_3 = -0.951411$, $\rho_4 = 0.459094$

SOLE OUTSIDE ROOT is HYPERBOLIC

CASE 5

• 2.0721961E-02

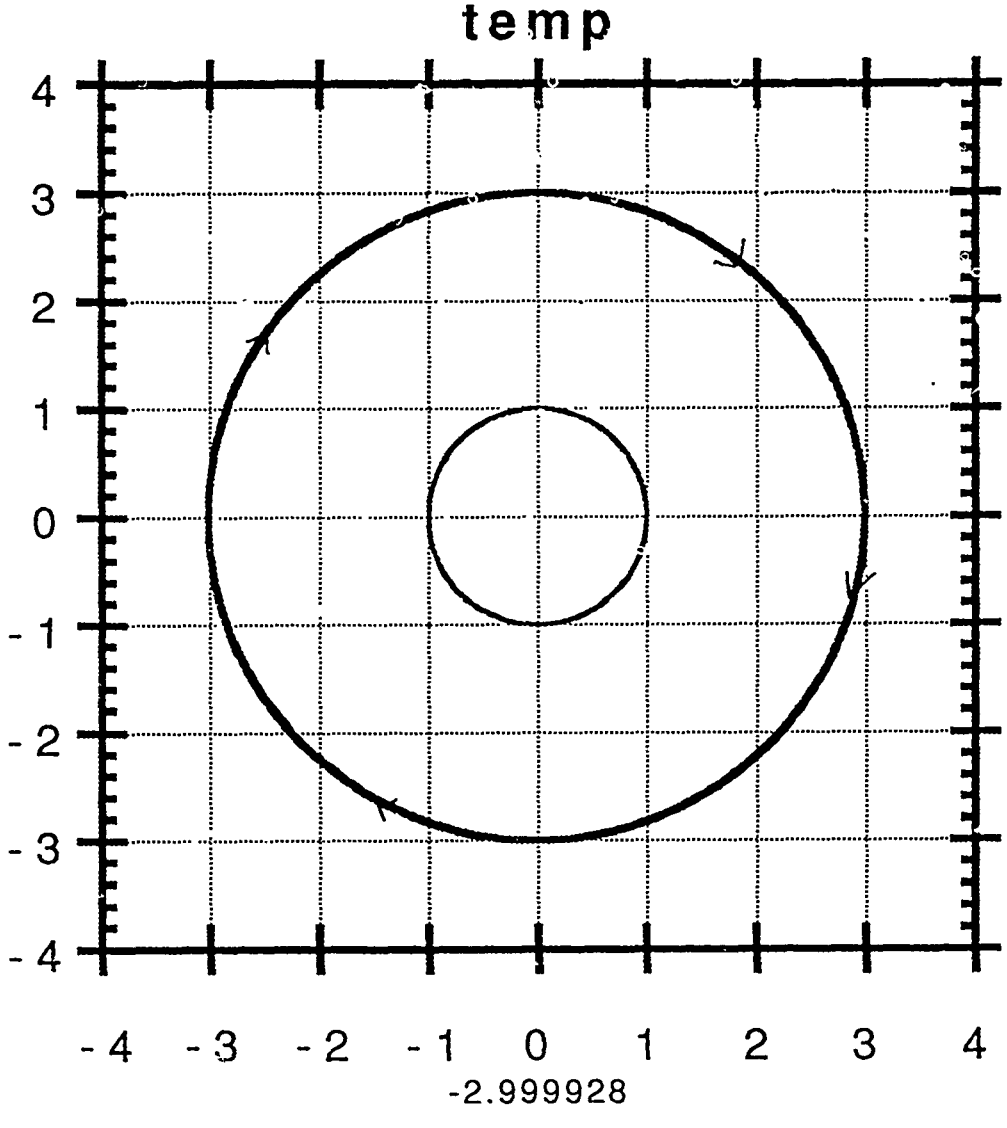

CASE 6

• 1.0000229E-03

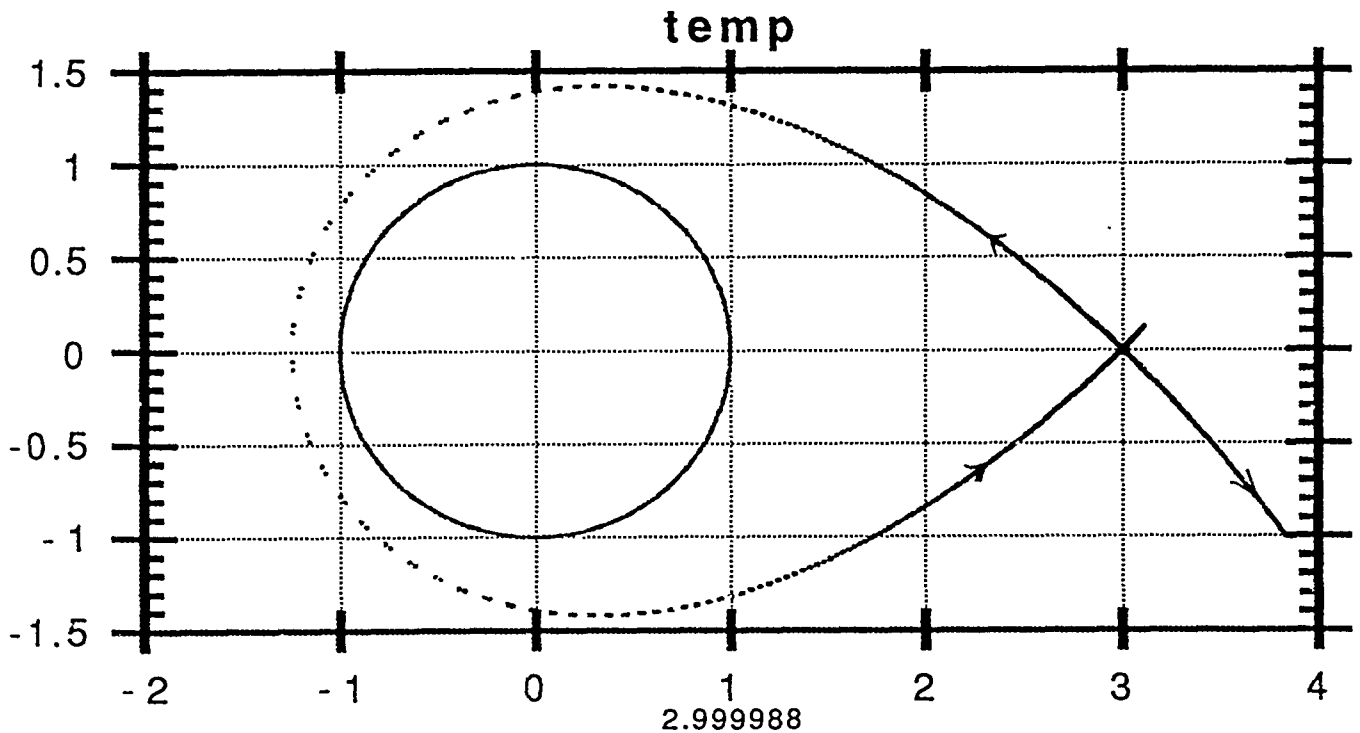

CASE 7 : $\sigma = -0.6321838$, $\sigma_0 = 2.62213$

$\rho_1 = 3.0$, $\rho_3 = -0.937964$, $\rho_2 = \rho_4 = 0.596137$

TWO ROOTS HAVE MERGED TO ONE, INSIDE BOUNDARY.

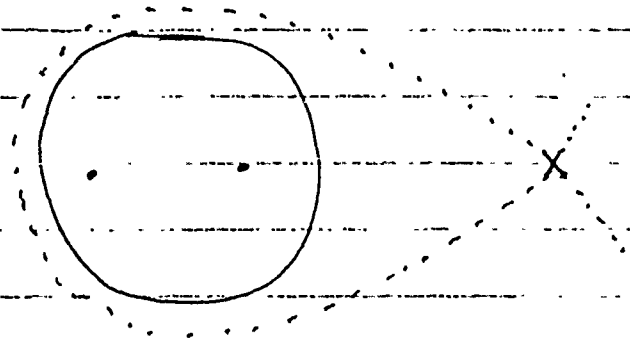

SEE PLOT

SOLE ROOT OUTSIDE BOUNDARY IS HYPERBOLIC.

CASE 8 : $-0.6321838 > \sigma > -98.518$

NOW ρ_2 AND ρ_4 COMPLEX, WITH $\rho_2 = \rho_4^*$

EXAMPLE : $\sigma = -10.0$, $\sigma_0 = -7.91667$

$\rho_1 = 3.0$, $\rho_3 = -0.125761$, $\rho_2 = -0.39545 + i\,1.57729$

$\rho_4 = \rho_2^*$

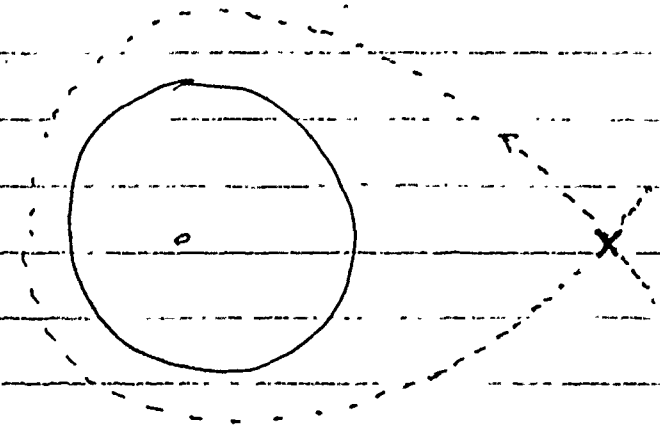

SOLE OUTSIDE POINT REMAINS HYPERBOLIC

CASE 7

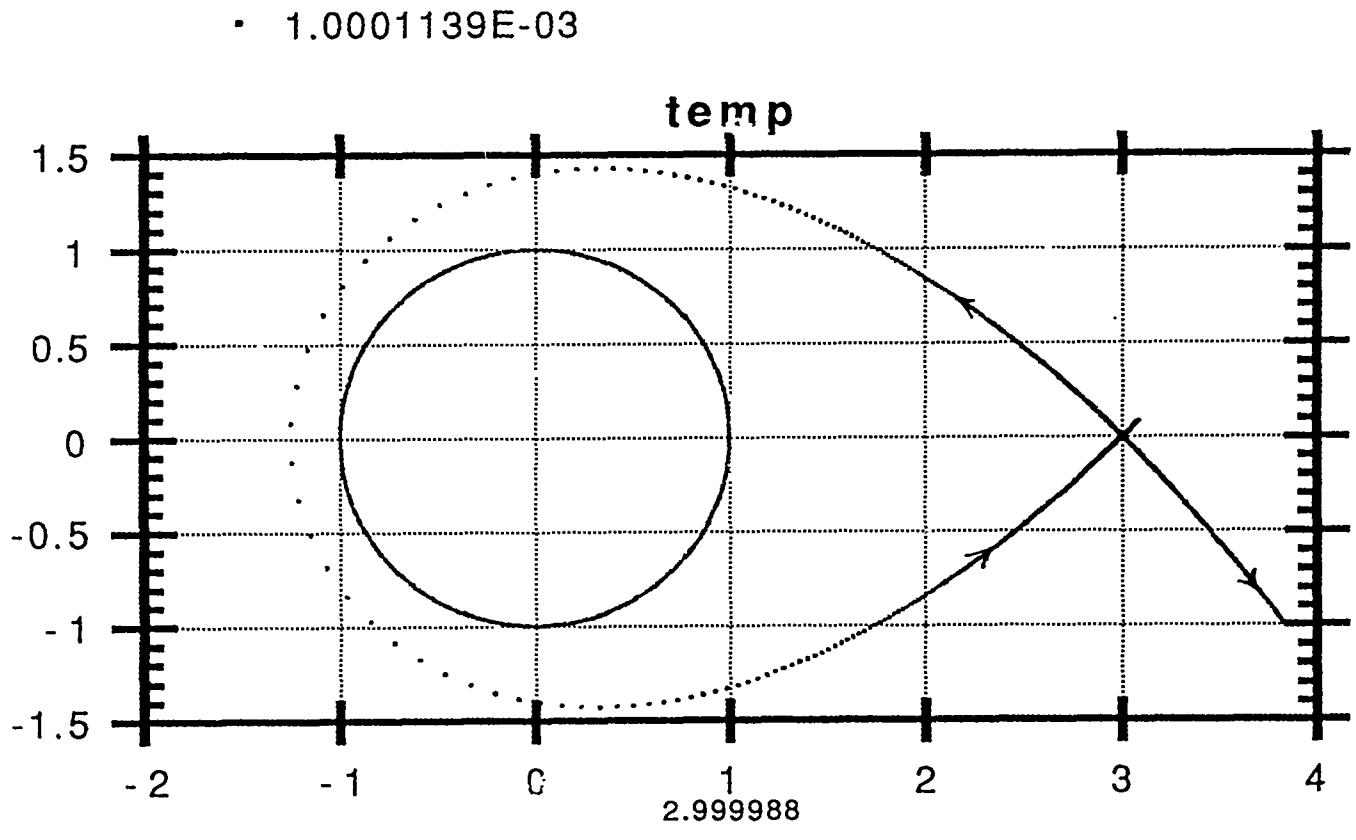

CASE 8

• 1.0000911E-03

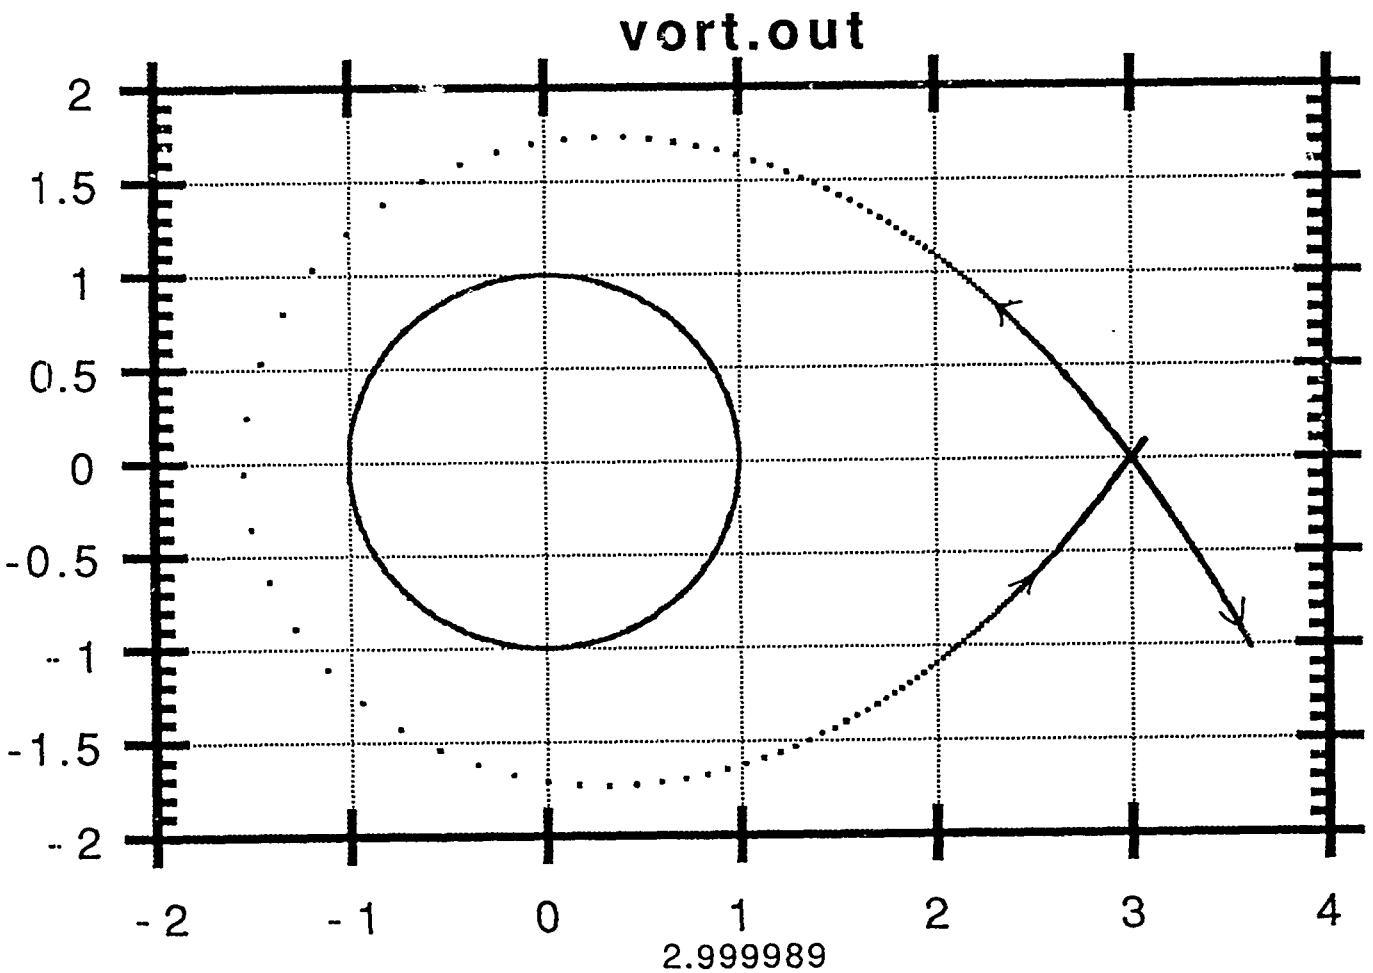

CASE 9: $\sigma = -98.518$ $\sigma_0 = -107.499$

ρ_2, ρ_4 BECOME REAL AGAIN, WITH

$\rho_1 = 3.0$, $\rho_3 = -0.009309$, $\rho_2 = \rho_4 = -6.0$

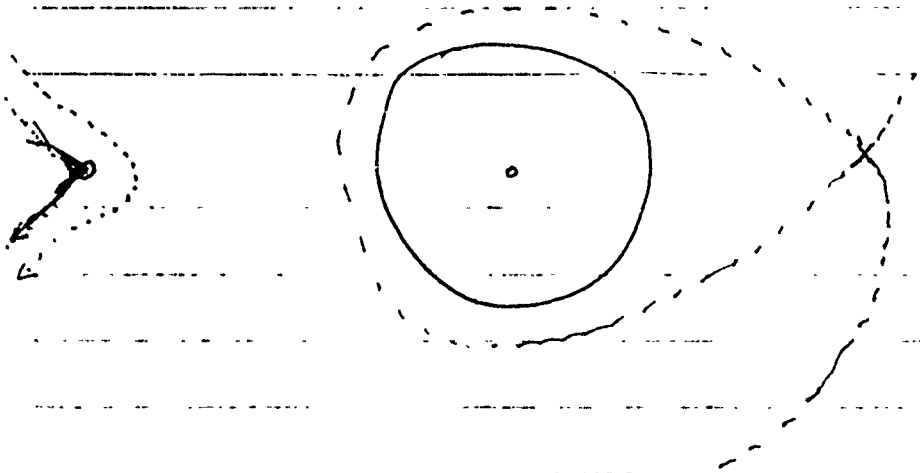

$\rho = -6.0$ IS STAGNATION POINT WITH ONLY ONE UNSTABLE DIRECTION COMING OUT, SIMILAR TO CASE 3.

CASE 10: $\sigma < -98.518$

EXAMPLE

$\sigma = -110.0$, $\sigma_0 = -120.417$

$\rho_1 = 3.0$, $\rho_2 = -8.89676$ $\rho_4 = -4.5116$, $\rho_3 = -0.008304$

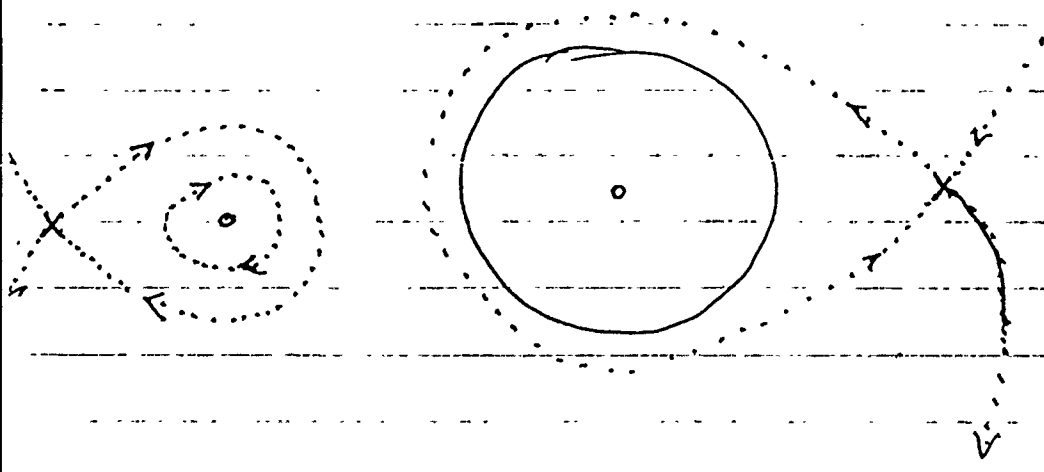

AGAIN HAVE 2 HYPERBOLIC AND 2 ELLIPTIC

CASE 9

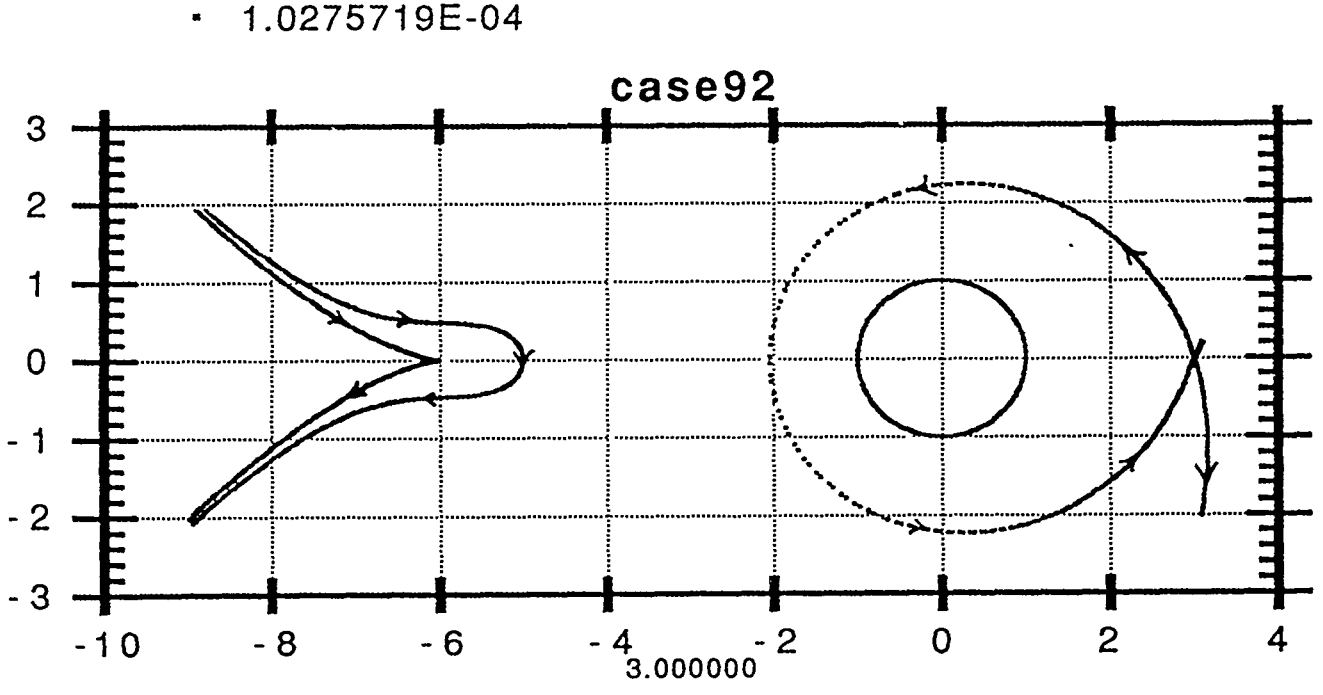

CASE 10

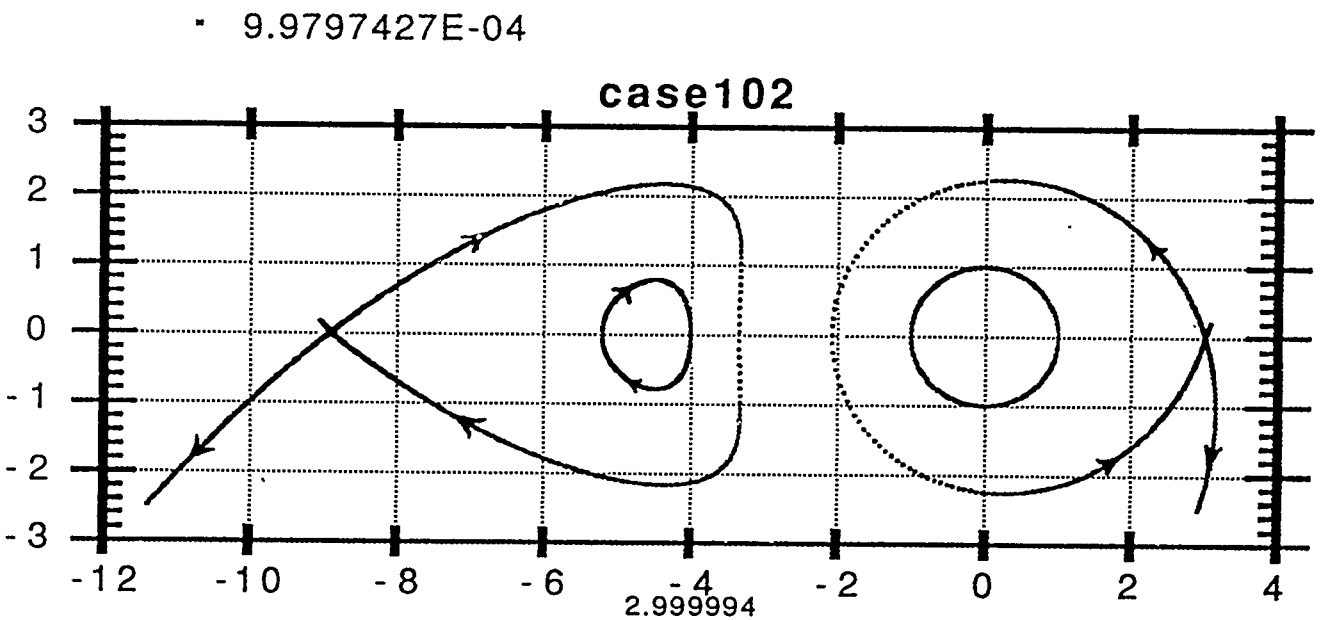

CASE. II $\sigma = -\infty$ $\sigma_0 = -\infty$

$\rho_1 = 3.0$, $\rho_2 = -\infty$, $\rho_3 = 0.0$, $\rho_4 = -3.0$

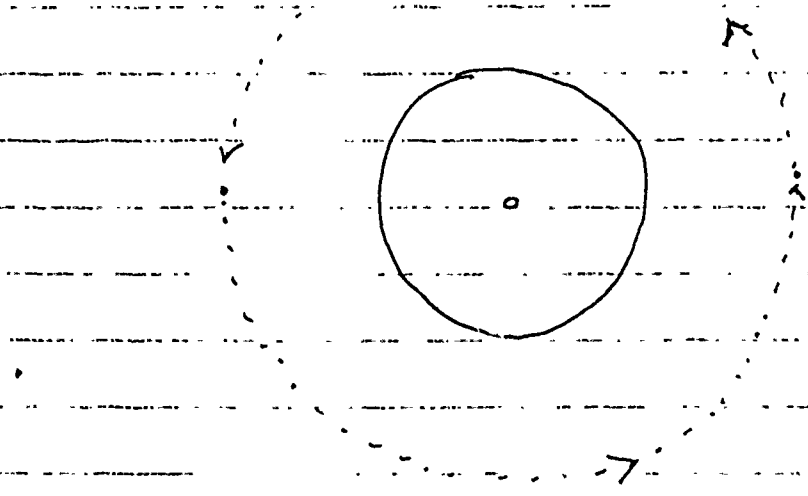

SIMPLE ROTATION THROUGH POINTS

NO HYPERBOLIC OR ELLIPTIC CHARACTER.

CASE 11

• 0.1035907

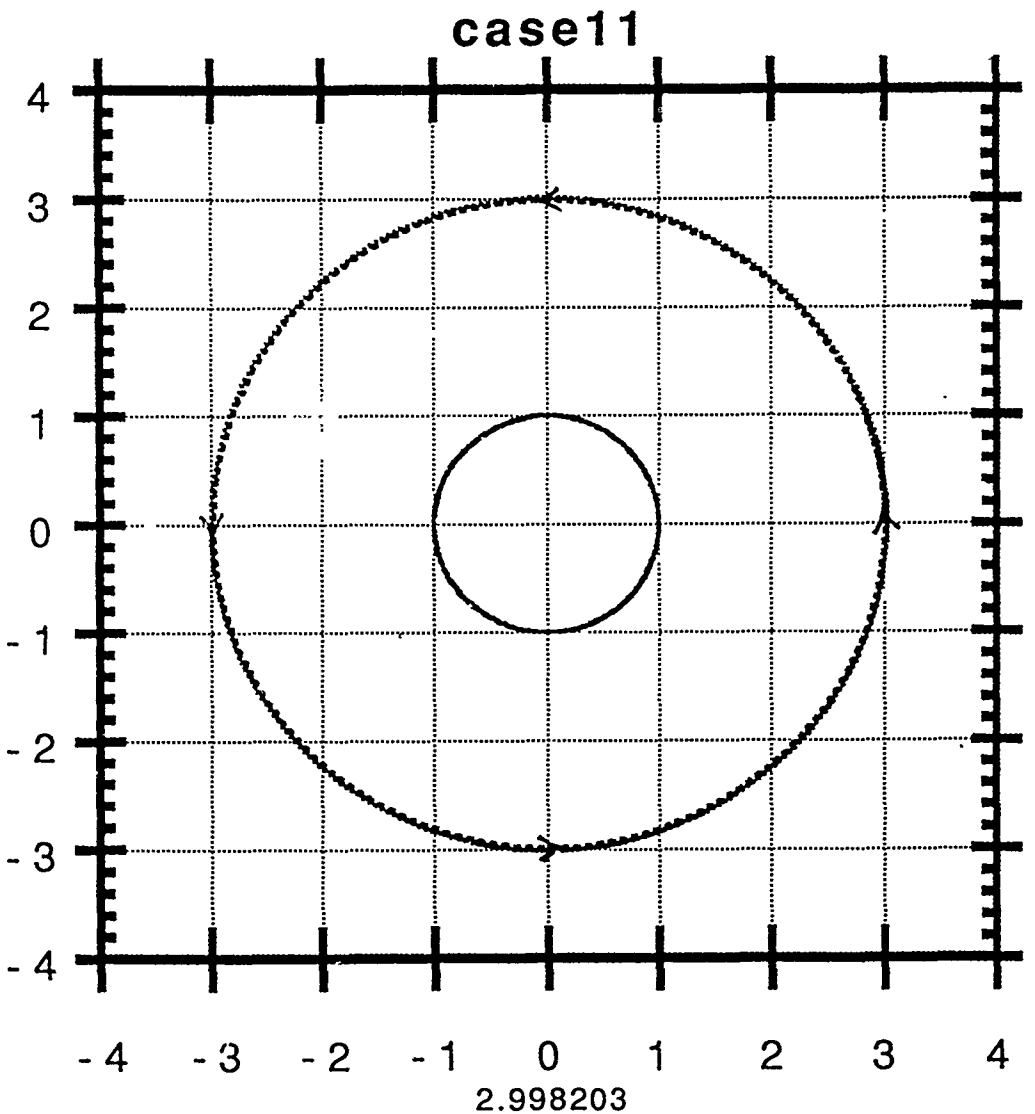

4 Revolution Capture trajectory showing Separatrix

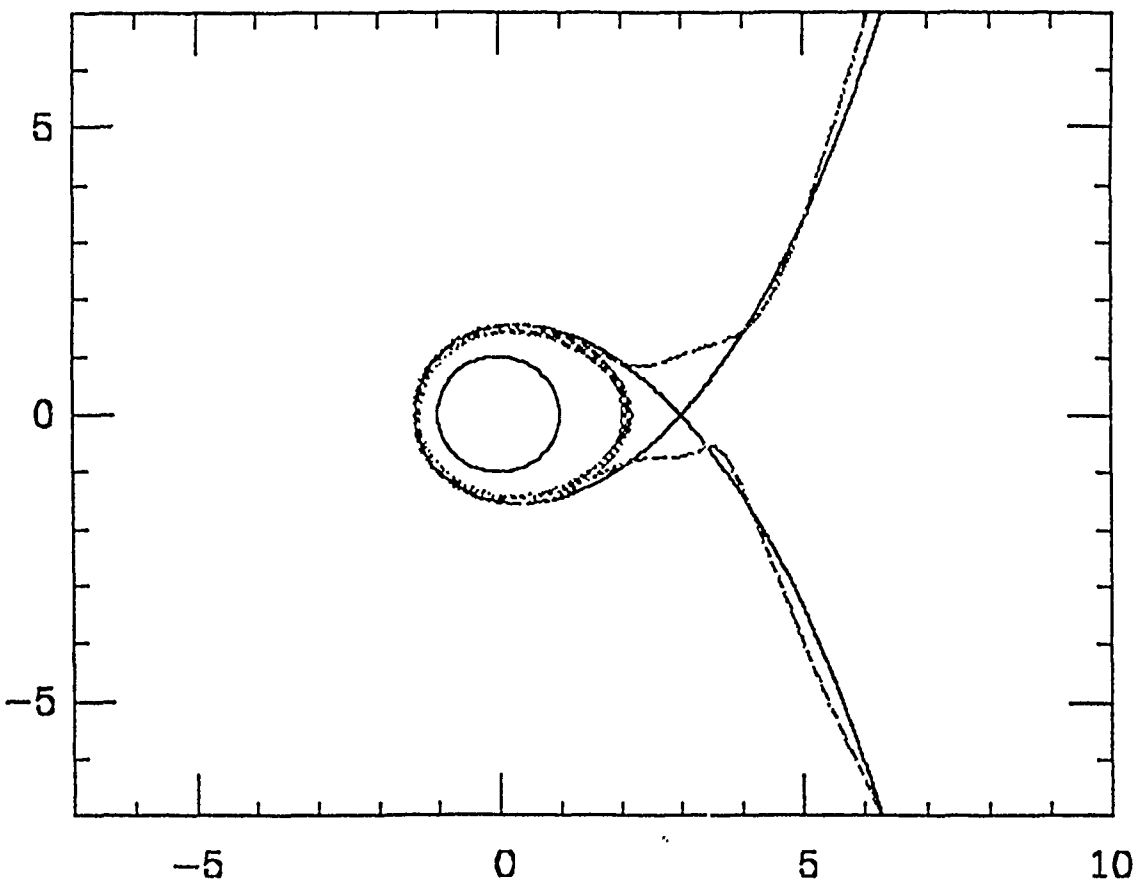

18 Revolution Capture

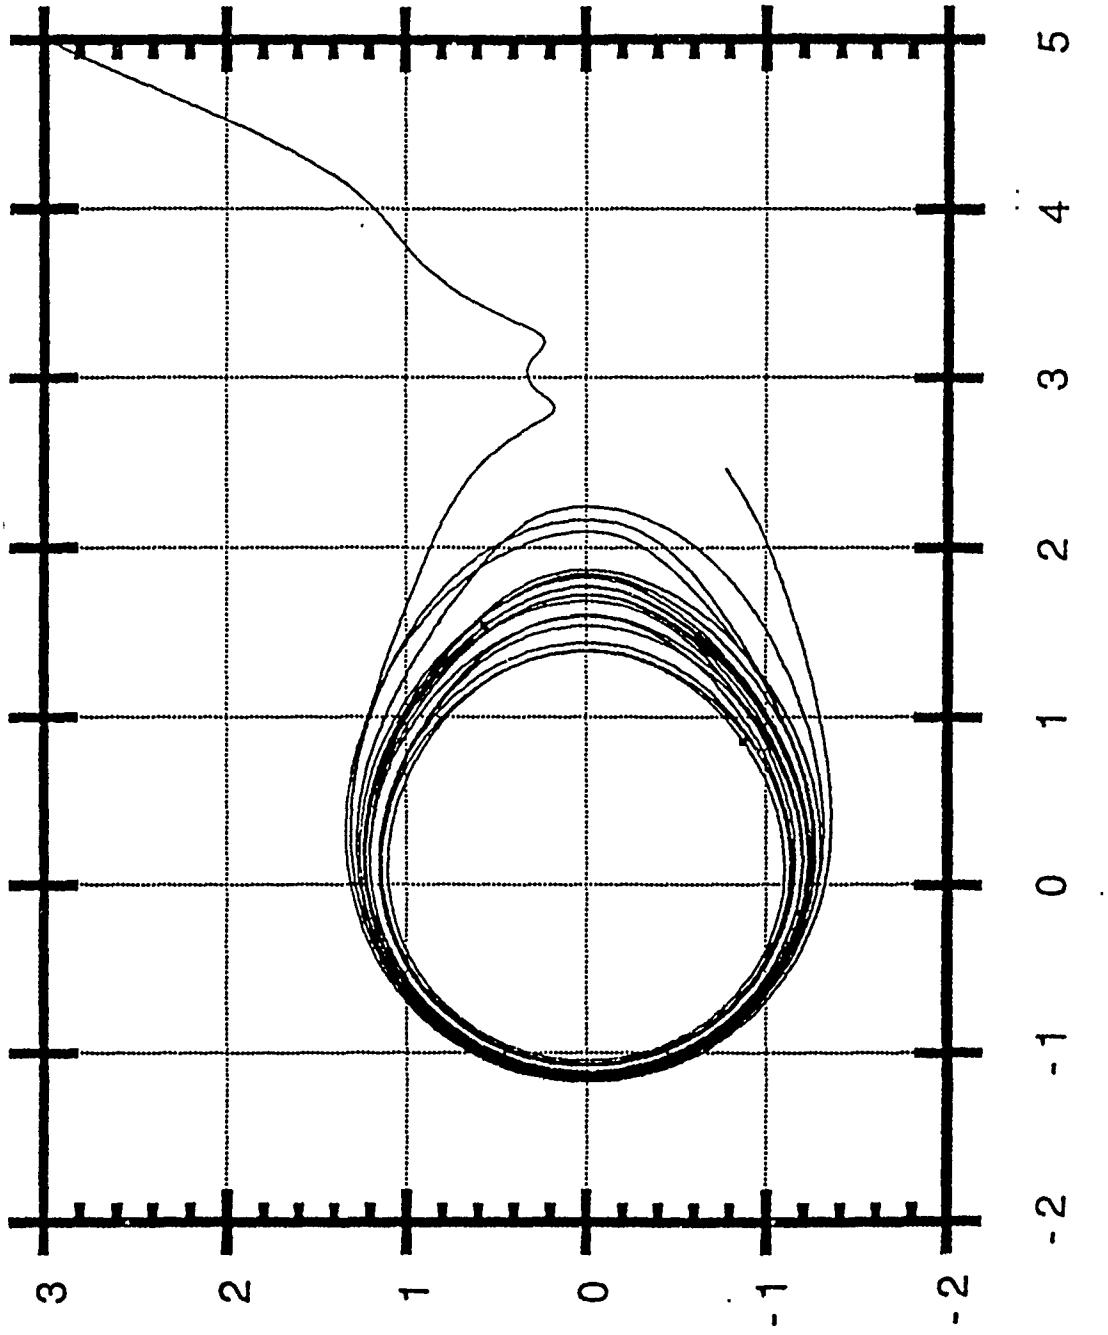

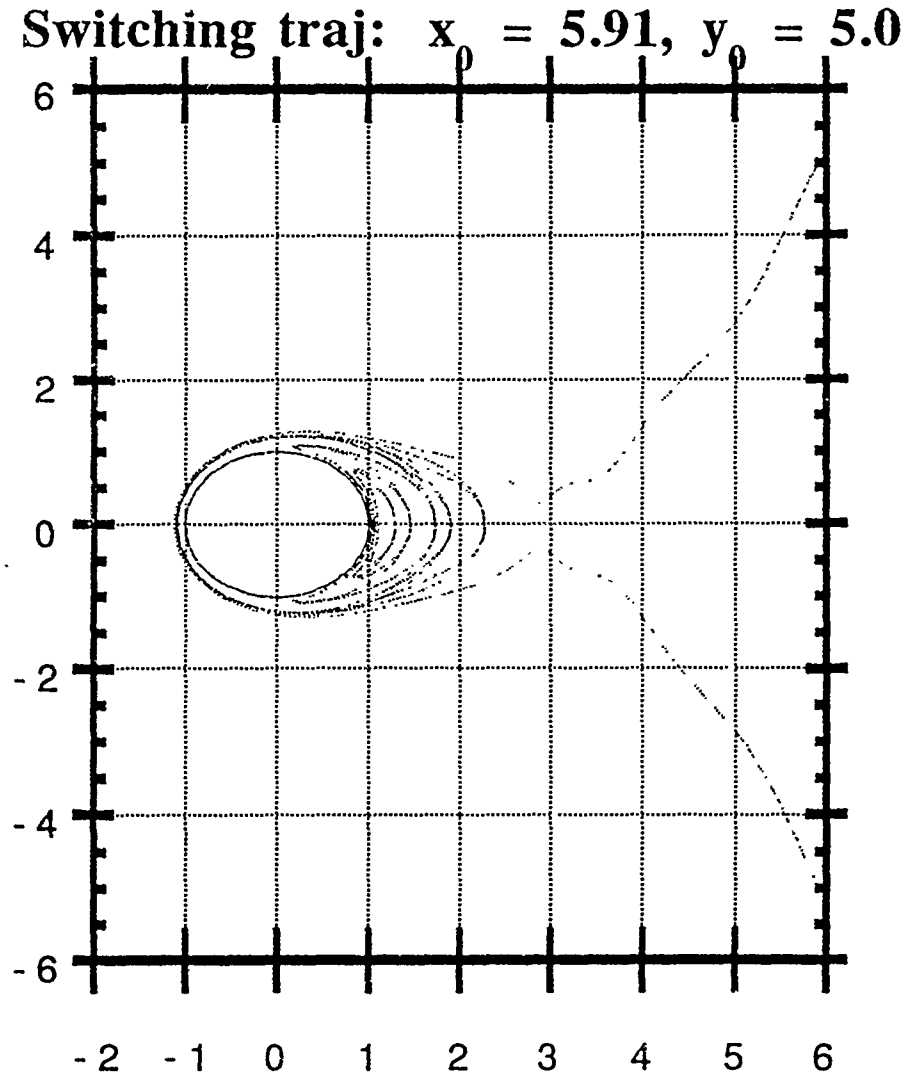
Switching traj: x_0 = 5.91, y_0 = 5.0
6
4
2
0
-2
-4
-6
-2
-1
0
1
2
3
4
5
6

Poincare Section Around Elliptic Point

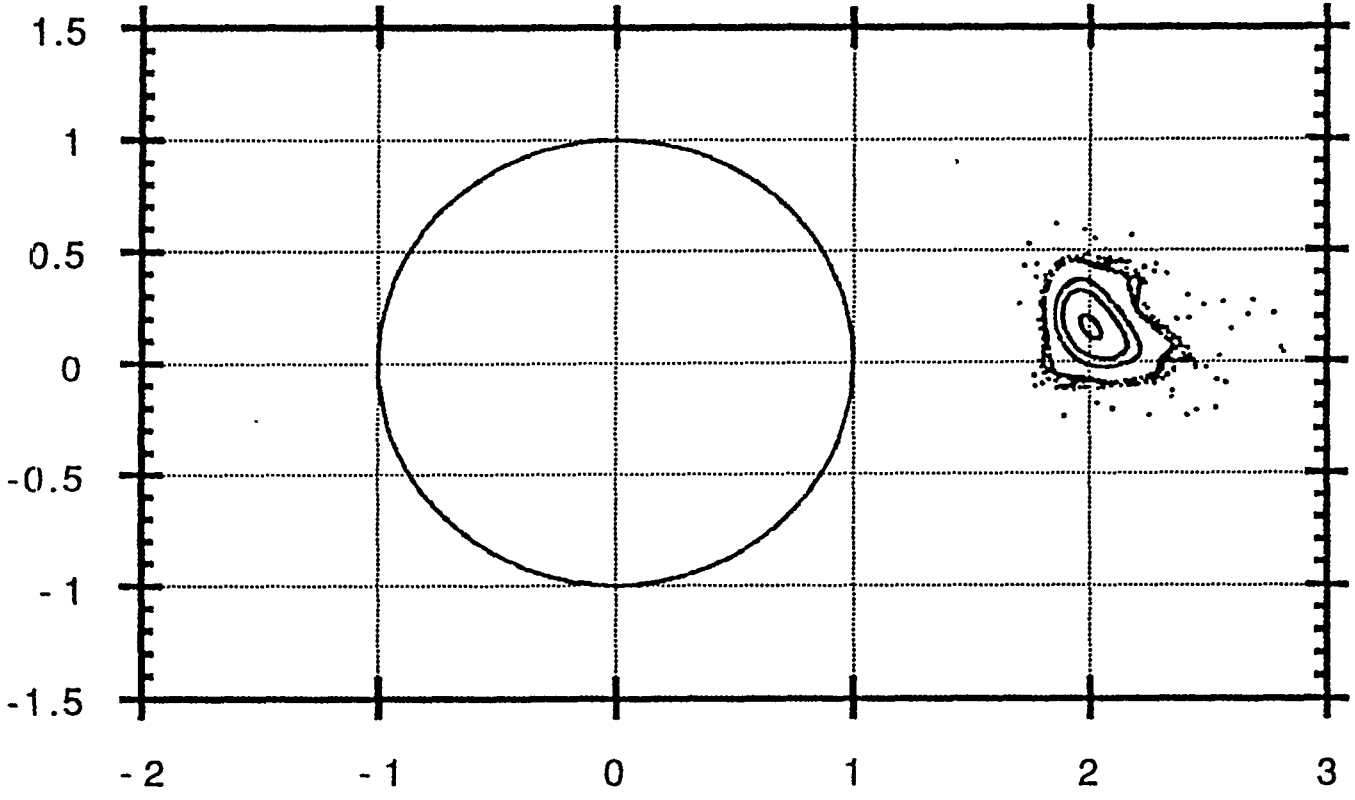

APPENDIX C

Photograph Of Chaotic Vortex Interaction Region Of Flow

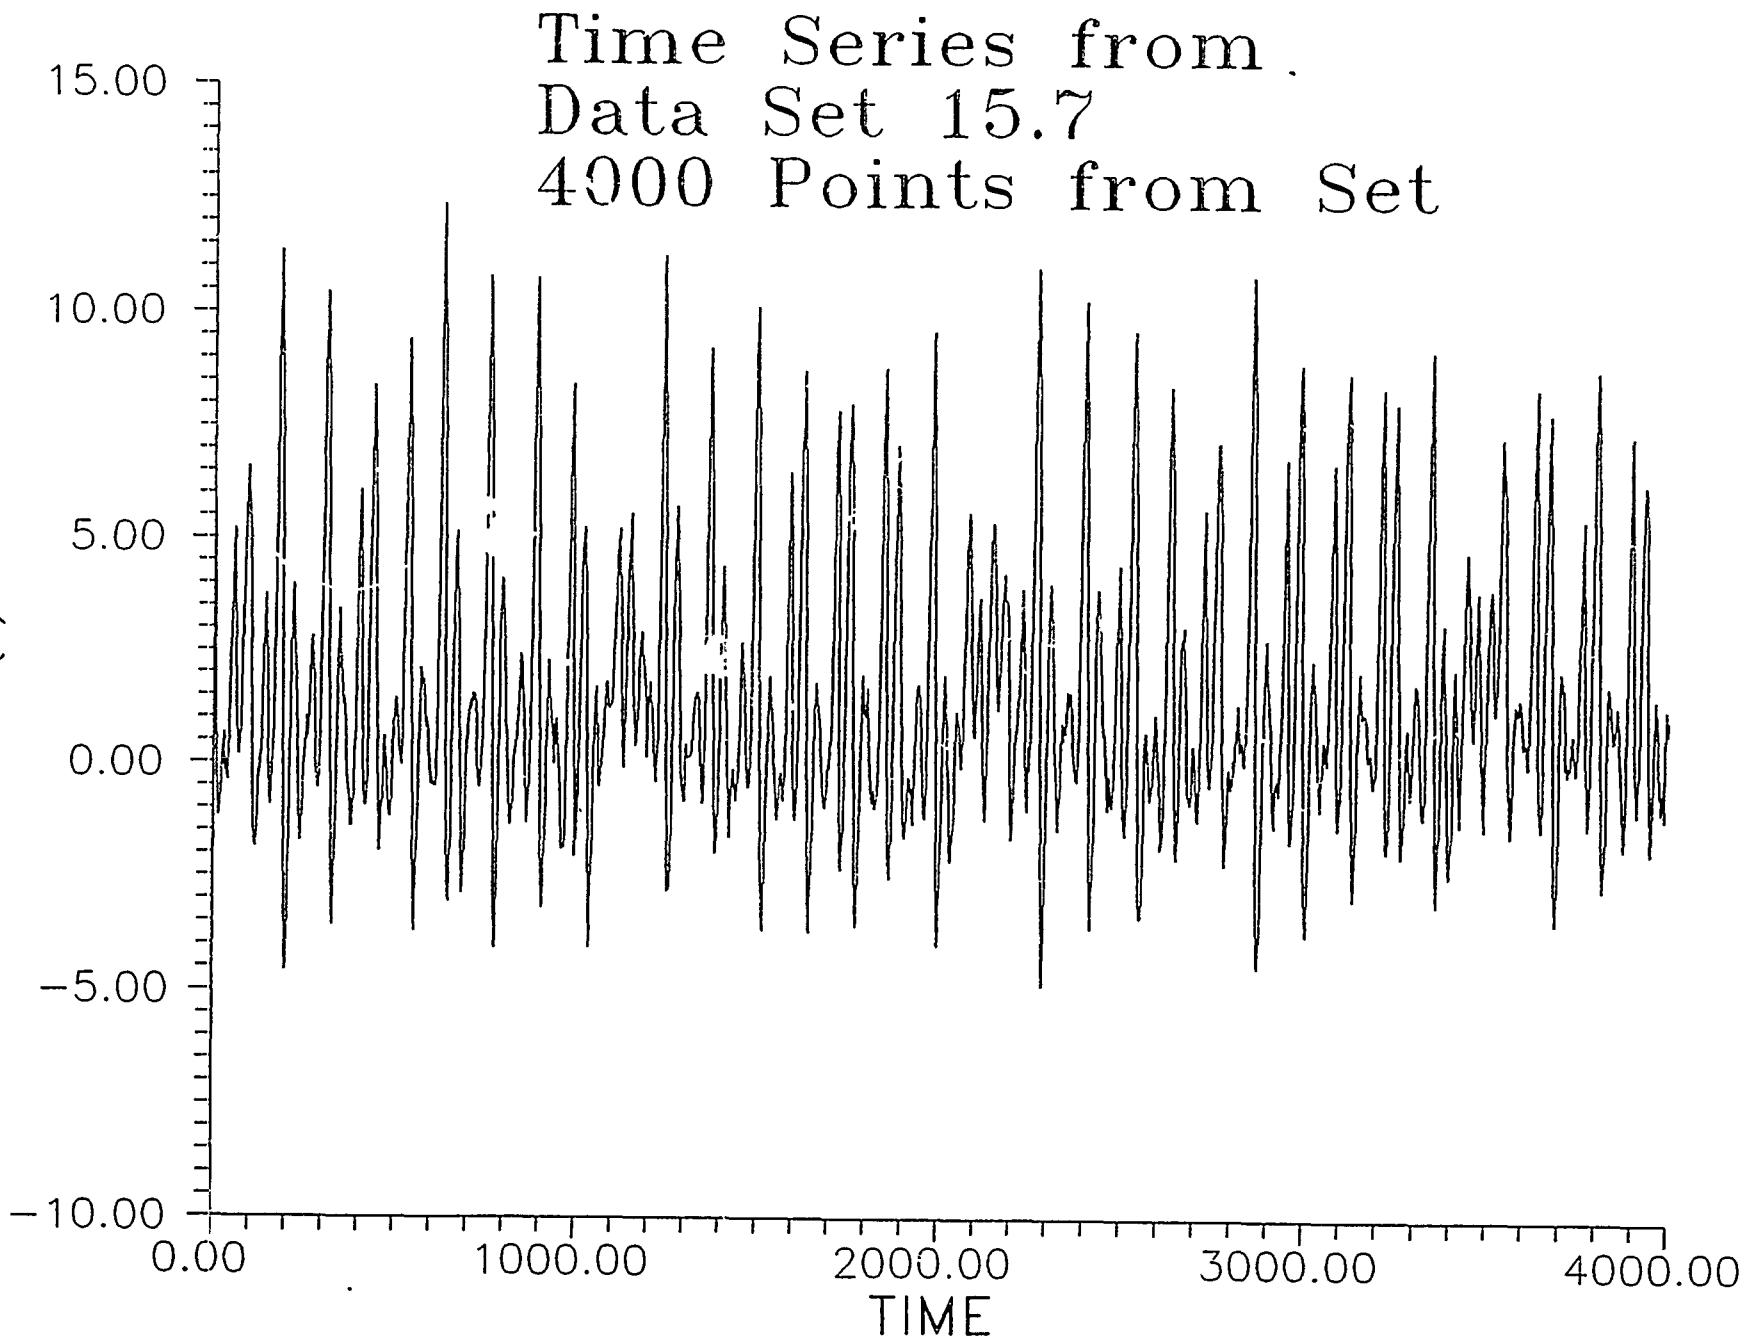
Time Series from
Data Set 15.7
4000 Points from Set
15.00
10.00
5.00
0.00
−5.00
−10.00
S(t)
0.00
1000.00
2000.00
3000.00
4000.00
TIME

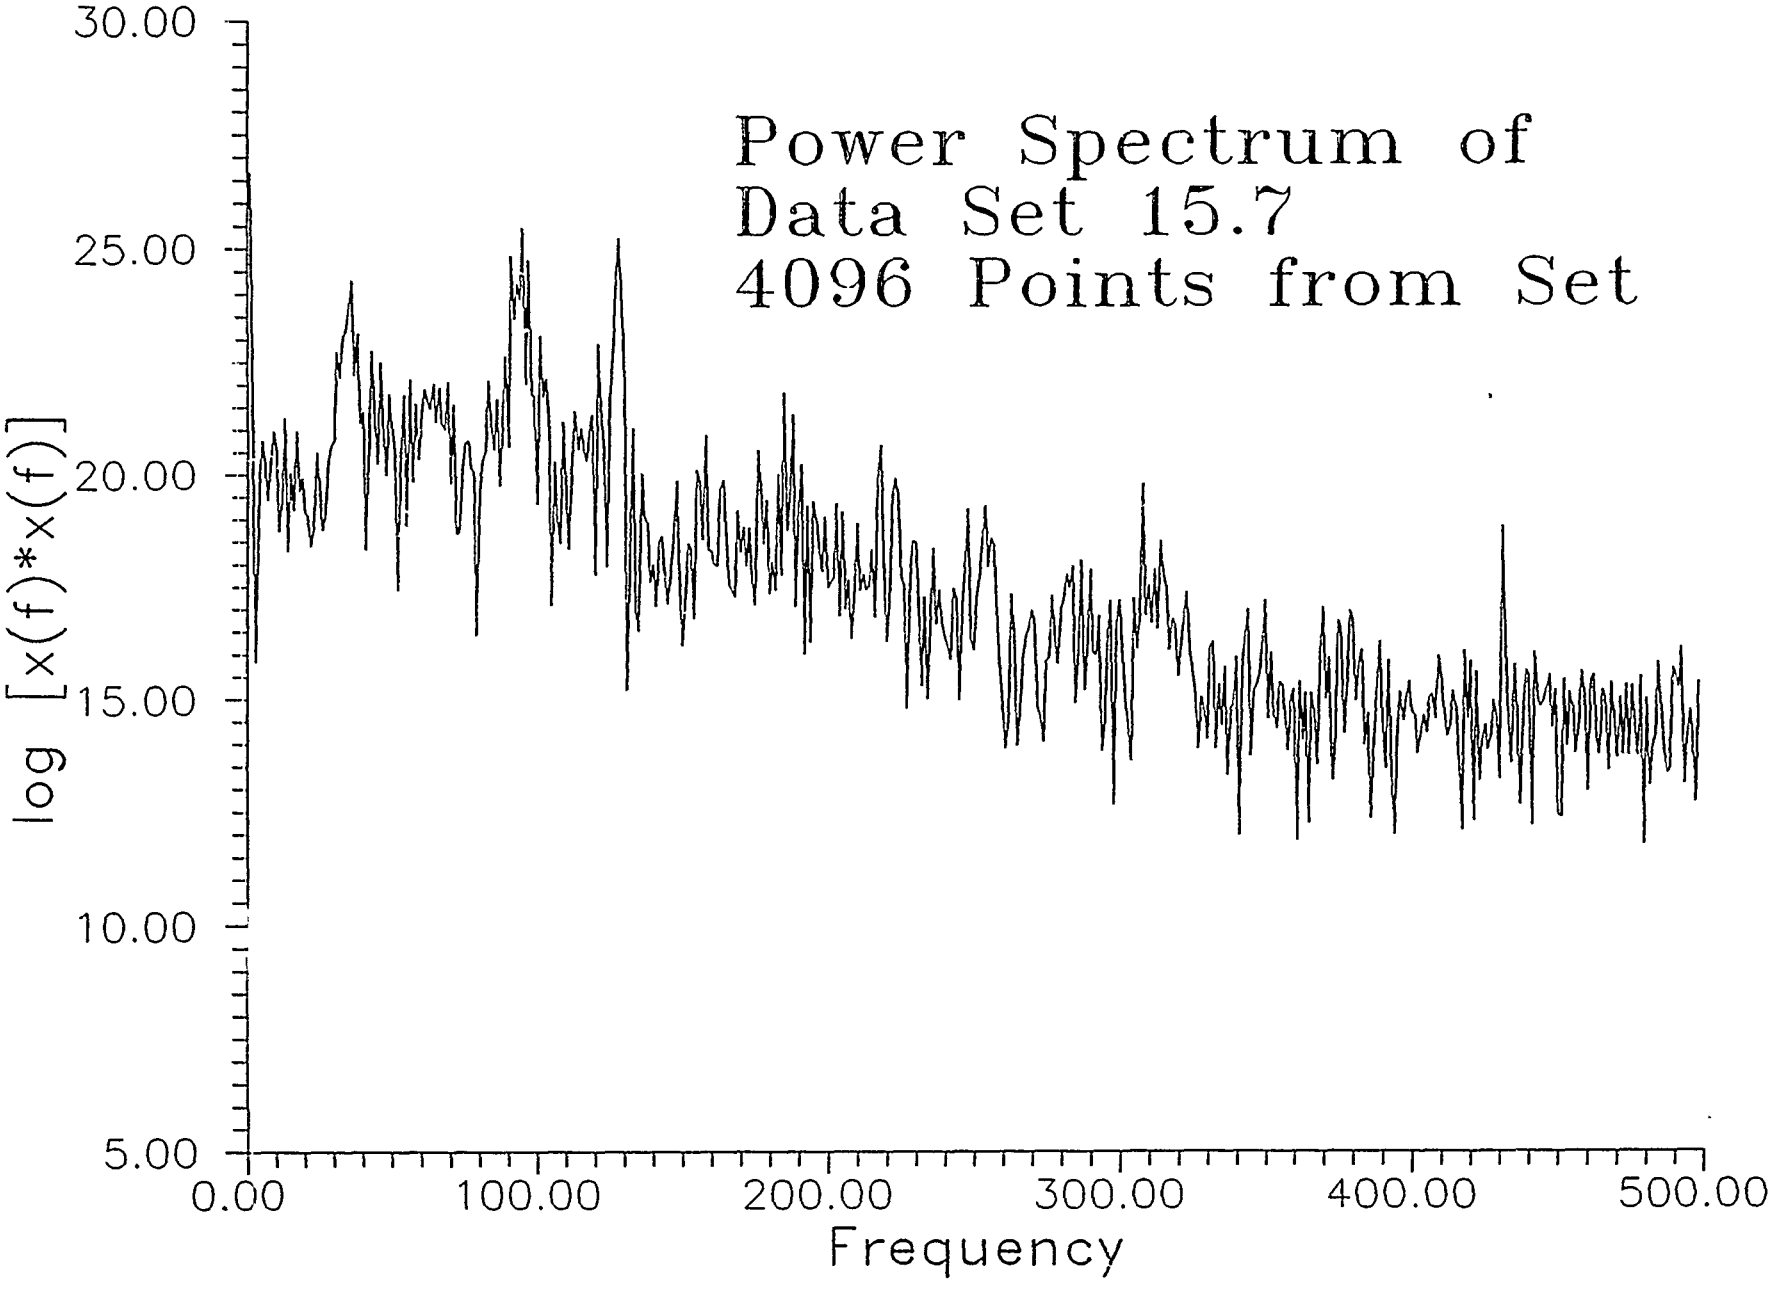
Power Spectrum of
Data Set 15.7
4096 Points from Set
log [x(f)*x(f)]
30.00
25.00
20.00
15.00
10.00
5.00
0.00
100.00
200.00
300.00
400.00
500.00
Frequency

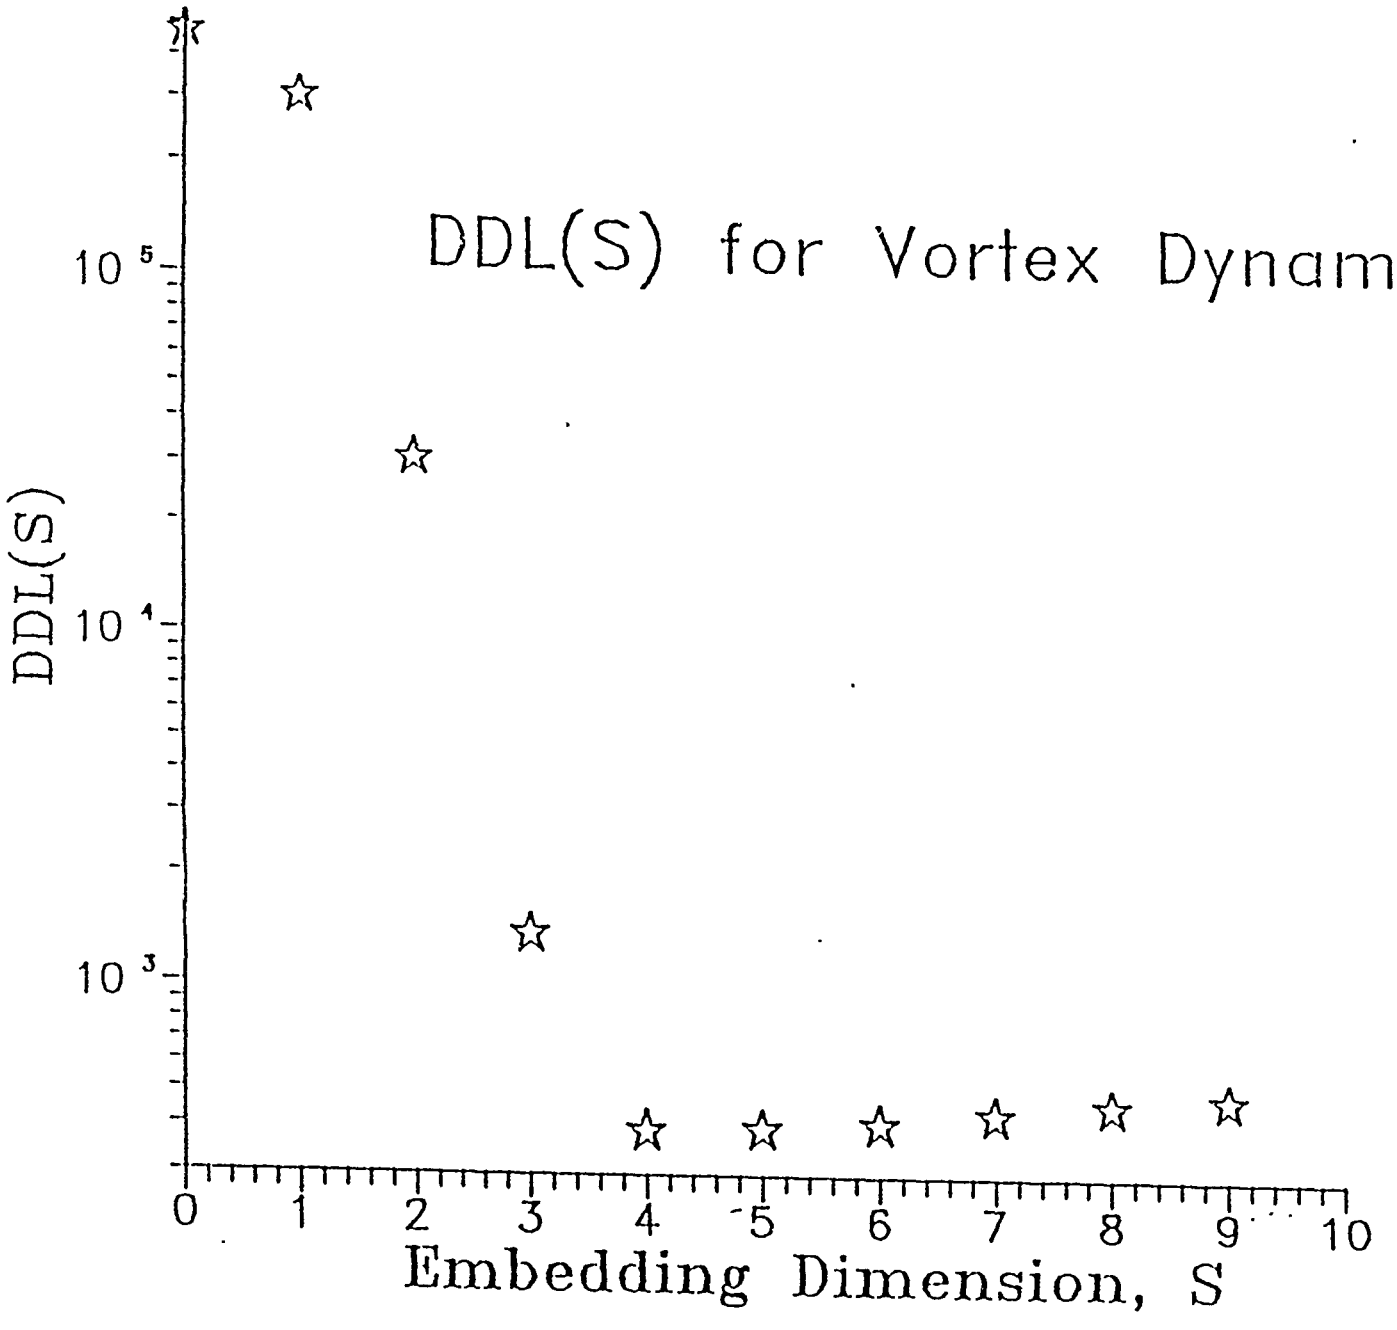
DDL(S) for Vortex Dynamics
DDL(S)
10 5
10 4
10 3
0 1 2 3 4 5 6 7 8 9 10
Embedding Dimension, S

FIGURE 3

APPENDIX D

PREDICTION AND SYSTEM IDENTIFICATION IN CHAOTIC NONLINEAR SYSTEMS: TIME SERIES WITH BROADBAND SPECTRA

Henry D.I. ABARBANEL [1]

Marine Physical Laboratory, Scripps Institution of Oceanography and Department of Physics, University of California, San Diego, Mail Code R-002, La Jolla, CA 92093-0402, USA

Reggie BROWN and James B. KADTKE

Institute for Nonlinear Science, University of California, San Diego, Mail Code R 002, La Jolla, CA 92093-0402, USA

Received 19 January 1989; revised manuscript received 24 March 1989, accepted for publication 18 April 1989
Communicated by D.D. Holm

We consider the problem of prediction and system identification for time series having broadband power spectra which arise from the intrinsic nonlinear dynamics of the system. We view the motion of the system in a reconstructed phase space which captures the attractor (usually strange) on which the system evolves, and give a procedure for constructing parameterized maps which evolve points in the phase space into the future. The predictor of future points in the phase space is a combination of operation on past points by the map and its iterates. Thus the map is regarded as a dynamical system, not just a fit to the data. The invariants of the dynamical system – the Lyapunov exponents and aspects of the invariant density on the attractor – are used as constraints on the choice of mapping parameters. The parameter values are chosen through a least-squares optimization procedure The method is applied to "data" from the Hénon map and shown to be feasible. It is found that the parameter values which minimize the least-squares criterion do not, in general, reproduce the invariants of the dynamical system. The maps which do reproduce the values of the invariants are not optimum in the least-squares sense, yet still are excellent predictors. We discuss several technical and general problems associated with prediction and system identification on strange attractors. In particular, we consider the matter of the evolution of points that are off the attractor (where little or no data is available), onto the attractor, where long-term motion takes place.

A broadband power spectrum observed in the time series of a system variable may have its origin in noise *extrinsic* to the system. However, it has become clear in recent years that its origin may be in the nonperiodic, deterministic chaos associated with a nonlinear system evolving on a finite-dimensional strange attractor [1]. In the latter case, which is the one we address in this note, the geometrical structure of the attractor and thus of the time series may be exposed by the method of phase space reconstruction [2–4]. This takes an observed scalar variable, $x(n) = x(t_0 + n\Delta t)$, and produces a D-dimensional embedding space from the time lagged signals $x(n)$, $x(n+\tau_1)$, ..., $x(n+\tau_{D-1})$, where the τ_i are appropriately chosen lags [5]. The sequence of D-vectors for $n = 1, 2, ..., N$

$$y(n) = [x(n), x(n+\tau_1), ..., x(n+\tau_{D-1})] \quad (1)$$

describes the evolution of the system in the embedding space. Theorems due to Takens and Mañé [3,4] tell us that if D is about twice the Hausdorff dimension of the attractor, we are assured of a good representation of that attractor by the $y(n)$. In practice, and in this note, a good representation is achieved by taking $\tau_j = j\tau$, with τ a common lag, and D the least integer dimension greater than the Hausdorff dimension. Along with others, we use properties of a correlation function [6] to decide which D is appropriate for a given data set.

The evolution of vectors y in the embedding space R^D provides a dynamically sound setting for the analysis of the broadband time series. The first step in our analysis is system identification or parameter estimation for a map $F(y, a)$ from R^D to itself, de-

[1] Institute for Nonlinear Science.

0375-9601/89/$ 03.50 © Elsevier Science Publishers B.V.
(North-Holland Physics Publishing Division)

pending on parameters $\boldsymbol{a}=(a_1, \ldots, a_p)$. This map takes $y(n)$ to $y(n+1)$:

$$y(n+1)=F(y(n), \boldsymbol{a}) . \tag{2}$$

Given a form for $F(y, \boldsymbol{a})$, the parameters $\boldsymbol{a}$ are to be estimated. Both the form of $F(y, \boldsymbol{a})$ and the criteria for choosing $\boldsymbol{a}$ are the issues of this note. Using $F(y, \boldsymbol{a})$ for prediction of the future evolution of the underlying system or for its control will be the subject of subsequent papers [7].

Actual extrinsic noise complicates the analysis of time series and, of course, is inevitably present in any interesting time series. In this note we ignore the role of such extrinsic noise. We wish to separate the treatment of data with *intrinsic* broadband time series from the study of such data contaminated by *extrinsic* noise. We will return to the analysis of extrinsically contaminated chaotic motion in our later work [7].

One of our central assumptions is that the points $y(n)$ lie on an attractor which is usually strange or fractional dimensional. Motion on the attractor is taken to be chaotic or sensitive to initial conditions. Thus, predicting the value of points, $y(n)$, on any individual orbit evolving from a starting value $y(1)$:

$$\begin{aligned} y(n)&=F(y(n-1), \boldsymbol{a})=F(F(y(n-2), \boldsymbol{a}), \boldsymbol{a}) \\ &=F^2(y(n-2), \boldsymbol{a})=\ldots=F^{n-1}(y(1), \boldsymbol{a}) \end{aligned} \tag{3}$$

is a numerically uncertain matter as n grows large. Indeed, for familiar low-dimensional systems such as the Hénon map or the Lorenz attractor [8], n larger than order 10 is usually unpredictable given small machine or initial condition errors.

Prediction of the evolution of a point, y, on the attractor in the reconstructed phase space may be determined by looking at points in the temporal past of y, *as well as* the evolution of points that are both spatially nearby and on the attractor. Knowledge of where the neighbors of the point y have evolved is as important as knowledge of where its temporal predecessors have been. This notion is present in several papers dealing with the same general subject as this note [9]. Our formulation of the concept is embodied in the explicit analytic formula we give for the map $F(y, \boldsymbol{a})$. This aids in both the analysis of the map's properties and its use in numerical work. With the idea in mind of using information about the neighbors of a point to assist in predicting where it will go, we choose mapping functions $F(y, \boldsymbol{a})$ of the form[#1]

$$F(y, \boldsymbol{a})=\sum_{j=1}^{N-1} y(j+1) f_\sigma(y, y(j), \boldsymbol{a}) , \tag{4}$$

where the $y(j)$ are the D-vectors constructed from the original scalar time series $x(n)$. The $f_\sigma(y, y(j), \boldsymbol{a})$ are functions parameterized by the $\boldsymbol{a}$'s and vanish rapidly when $|y-y(j)|^2/\sigma \gg 1$. ($|\ |$ is the Euclidean distance in $\mathbb{R}^D$.) $F(y(n), \boldsymbol{a})$ is then determined by all the neighbors of $y(n)$ regardless of the temporal order in which the neighborhood is visited. For appropriate choices of f_σ the value of $F(y, \boldsymbol{a})$ is a weighted average of the places in $\mathbb{R}^D$ to which the neighbors of y go in one iteration of the map.

Our specific choice of $f_\sigma(y, y(j), \boldsymbol{a})$ is

$$\begin{aligned} f_\sigma(y, y(j), \boldsymbol{a})&=\exp(-|y-y(j)|^2/\sigma) \\ &\times\Big(a_1+a_2 y(j)\cdot[y-y(j)] \\ &+\sum_{k=3}^{P} a_k(|y-y(j)|^2/\sigma)^{m_k}\Big) \end{aligned} \tag{5}$$

with m_k some choice of integers. This is one among a large class of f's satisfying our requirements of evolving points according to the fate of their neighbors. This equation determines how "close" y is to each of the data vectors $y(j)$. For large values of $|y-y(j)|^2/\sigma$ the exponential dominates and the function approaches zero. The polynomial softens the exponential decay while simultaneously providing a polynomial fit to the data for those values of $|y-y(j)|^2/\sigma$ that are small.

A problem arises when $f_\sigma(y, y(j), \boldsymbol{a})$ in our map depends only on the Euclidean distance of $y-y(j)$. We find that if σ is small enough to accurately forecast points on the attractor then the Jacobian matrix of $F(y, \boldsymbol{a})$ which we will numerically evaluate on the data may be so small (due to the exponential term) that we are unable to calculate Lyapunov exponents. The term associated with a_2 provides for a non-zero Jacobian when y is evaluated at some $y(j)$ which has no neighbors in the data set. A zero Jacobian would

[#1] Arguments for the general form here are given in ref. [10].

be fatal, since we will eventually use $F(y, a)$ to calculate our Lyapunov exponents.

The predictor we choose for a point $y(N+1)$ could in principle be just $F(y(N), a)$. It could also be $F^2(y(N-1), a)$ or $F^j(y(N-j+1), a)$ for any j. We are seeking a map, $F(y, a)$, to reproduce the data $y(1), ..., y(N)$ as well as possible, However, it is also critical that the value of $y(j)$ be the result of iterates of the map on $y(j-1), y(j-2), ...$ etc; that is, we are approximating the dynamical system which gives us the $y(j)$. With this in mind we have chosen as predictors a weighted average over L equivalent forms of $y(N+1)$. L will typically be a small number, and the weights X_j will be chosen to decrease as j increases in some manner consistent with one's confidence in the reliability of $F^j(y, a)$. We choose then

$$y(N+1)=\sum_{j=1}^{L} X_j F^j(y(N-j+1, a)\,. \tag{6}$$

If $F(y, a)$ were the exact mapping, then each term in this sum would be $X_j y(N+1)$ itself, thus we require

$$\sum_{j=1}^{L} X_j = 1\,. \tag{7}$$

This choice of predictor is a natural extension to the nonlinear situation of the linear predictive scheme

$$y(N+1)=\sum_{j=1}^{L} X_j y(N-j+1) \tag{8}$$

discussed in many places [11]. By including iterations of the map $F(y, a)$ it explicitly embodies the idea that we are dealing with an iterated map or dynamical system. It also provides a "lever arm" on predictions since it looks back not just one step but many to see where a given point y will evolve. Since $F^j(y, a)$ provides information on the evolution of the neighbors of the points which end up at $y(N+1)$ this "lever arm" is both temporal and spatial. Interestingly, using the optimization criteria we are about to discuss, this method also yields much better quantitative fits to the data than using the single term $F(y(N), a)$.

In principle one would establish the parameters $X=(X_1, ..., X_L)$ and $a=(a_1, ..., a_P)$ by minimizing the *cost function*

$$C(X, a)=(N-L-1)^{-1} \times \sum_{k=L}^{N-1} \left| y(k+1)-\sum_{j=1}^{L} X_j F^j(y(k-j+1), a) \right|^2 . \tag{9}$$

This alone would be a fairly standard least-squares way of determining the parameters X and a. Other than the requirement that $F(y, a)$ search phase space neighborhoods to determine where to map y, this minimization is familiar. In this paper we report results obtained by minimizing $C(X, a)$ by searching various values of a, for *fixed* values of X. In our subsequent paper we discuss searches over both X and a.

The individual orbits of dynamical systems of the form $y \to F(y, a)$ which exhibit chaotic motion are sensitive to changes in initial conditions or roundoff error in machine calculations [1]. There are, however, quantities that are *invariant* under the motion and are characteristic of the dynamical system that gives rise to the data, $y(n)$. The least-squares estimation of a by minimizing $C(X, a)$ does not guarantee that the resulting map $F(y, a)$ will give the correct invariants. Even if $C(X, a)=0$ we cannot *guarantee* that $F(y, a)$ captures the full dynamics of the system that generated the data, since all data are subject to error, we are using a finite data set, and only predicting a few steps into the future. In any event, in practice, we are able to make $C(X, a)$ small but nonzero. The fact that $C(X, a)$ is not zero implies that there is certainly no guarantee that the invariants will be reproduced by the map $F(y, a)$. To guarantee that the values of the invariants of the underlying dynamical system are built into the parameterized maps $F(y, a)$, we seek the parameters values that minimize $C(X, a)$ *subject to* the constraints that $F(y, a)$ yield the correct values for the invariant characteristics. To implement this we need to identify the relevant invariants, determine them from the data set in a manner independent of the least-squares minimization, and give rules on finding them for any map $F(y, a)$.

We are aware of two kinds of invariants for dynamical systems [1]. Both are connected with ergodic properties of the true underlying dynamics that generated the data set. The first is the set of characteristic Lyapunov exponents, $\lambda_1, \lambda_2, ..., \lambda_D$. The

second is the invariant density of points on the attractor, $\rho(y)$.

We assume that when $C(X, a)$ is near zero the map $F(y, a)$ is ergodic on the attractor. If this is true, then the λ_a^{map} are given by Oseledec's multiplicative ergodic theorem [12] as the logarithms of the eigenvalues of

$$[(DF^M(y(1)))^{\dagger} DF^M(y(1))]^{1/2M} \tag{10}$$

as $M \to \infty$, where the dagger denotes Hermitean conjugate, $DF(y)$ is the Jacobian matrix

$$DF(y)_{\alpha\beta} = \frac{\partial F_\alpha(y, a)}{\partial y_\beta} \tag{11}$$

and

$$DF^M(y) = DF(F^{M-1}(y, a)) \ldots DF(F(y, a)) DF(y) . \tag{12}$$

The λ's are invariant in the sense that all initial points $y(1)$ that are on the attractor yield the same values for the λ's.

If only λ_1^{map} is desired, then

$$\frac{1}{M} \log\{\mathrm{Tr}[DF^M(y)]\} \tag{13}$$

gives a very accurate value for λ_1^{map} as M becomes large. Basically this is because

$$\mathrm{Tr}[DF^M(y)] = \sum_{a=1}^{D} \exp(M\lambda_a) \approx \exp(M\lambda_1) \tag{14}$$

for large M. This rule for finding λ_1^{map} from the map $F(y, a)$ turns out, in practice, to be easy to program and completely adequate for use in constraining the optimization of $C(X, a)$.

Establishing the Lyapunov characteristic exponents from data turns out to be a delicate procedure for all but the largest positive exponent, λ_1^{data} [13]. This appears to be operationally the case whether the data comes from an experiment or is computer generated. Hence, we have chosen to restrict our attention here to the largest positive Lyapunov exponent. Let us suppose λ_1^{data} has been determined from the data, $y(n)$. The equality $\lambda_1^{\text{map}} = \lambda_1^{\text{data}}$ forms one of the constraints we will impose on the minimization of the cost function.

It seems to us a matter of some interest to create reliable, efficient methods for the determination of the full Lyapunov spectrum from data and from maps $F(y, a)$.

We now turn our attention to the second type of invariant, the invariant density, $\rho(y)$. The invariant density of a map, $F(y, a)$, is defined as

$$\rho(y)_{\text{map}} = \lim_{M \to \infty} \frac{1}{M} \sum_{k=1}^{M} \delta^D(y - F^k(y(1), a)) . \tag{15}$$

For an ergodic map (which we have assumed our map to be) $\rho(y)$ is invariant in the sense that all initial points $y(1)$ that are on the attractor yield the same value for $\rho(y)$. The spatial average of $g(y)$, an arbitrary function on phase space, is given by

$$\langle g \rangle = \int \mathrm{d}^D y \rho(y) g(y) .$$

For ergodic maps $\langle g \rangle$ is invariant under the action of $F(y, a)$, i.e.

$$\int \mathrm{d}^D y \rho(y)_{\text{map}} g(y) = \langle g \rangle_{\text{map}} = \int \mathrm{d}^D y \rho(y)_{\text{map}} g(F(y, a)) . \tag{16}$$

The invariant density determined by the data is clearly given as

$$\rho(y)_{\text{data}} = \frac{1}{N} \sum_{k=1}^{N} \delta^D(y - y(k)) . \tag{17}$$

Thus, $\langle g \rangle_{\text{data}}$ is given by

$$\langle g \rangle_{\text{data}} = \frac{1}{N} \sum_{k=1}^{N} g[y(k)] . \tag{18}$$

The second constraint we will impose on $C(X, a)$ is the equality of $\rho(y)_{\text{map}}$ with $\rho(y)_{\text{data}}$. We cannot constrain the minimization of $C(X, a)$ by the full content of $\rho(y)$, since it contains an infinite amount of local information. However, we can choose some specific functions $g_i(y)$ which we deem important about the dynamical system and require $\langle g_i \rangle_{\text{map}} = \langle g_i \rangle_{\text{data}}$ as constraints. For this paper we report the results obtained by using $g = |y|^4 \exp(-\frac{1}{20}|y|^2)$. There is no intrinsic significance to this phase function, but it does have contributions from large values of Euclidean distances on the attractor. Also it is not connected in any direct way with the function $C(X, a)$ and, thus, contains information about the attractor quite different from $C(X, a)$.

An obvious question is what is the best choice of moments, $\langle g \rangle$, to use to constrain the cost function.

To answer this we expand $\rho(y)$ in terms of some set of orthonormal functions $\psi_\mu(y)$ which are concentrated on the attractor, $\rho(y)=\sum_{\mu=1} B_\mu \psi_\mu(y)$. The requirement that the $\psi_\mu(y)$'s be concentrated on the attractor implies that the number of these functions needed to represent $\rho(y)$ accurately (within the resolution given by finite N) can be small. Thus the details of $\rho(y)$ can be transmitted in terms of a set of functions "tuned" to the attractor and their coefficients. One learns the functions from one data set and uses them as a basis set for future data sets. If we knew the B_μ, we would be able to reconstruct the phase space average of any $g(y)$ since

$$\langle g\rangle = \sum_{\mu=1} B_\mu \left(\int \mathrm{d}^D y\, \psi_\mu(y)\, g(y) \right). \tag{19}$$

The term in large parentheses is independent of the dynamical system. It depends on the phase function $g(y)$ and the basis vectors only.

The B_μ's are the moments (phase space averages) of the eigenfunctions $\psi_\mu(y)$ and constitute the optimal moments to use in constraining the cost function. In the larger paper which follows this note [7] we show how the $B_\mu = \int \mathrm{d}^D y\, \rho(y)\, \psi_\mu(y)$ may be used to constrain the minimization of $C(X, a)$. The key is the choice of the $\psi_\mu(y)$ which we select as the *optimum* Karhunen–Loève eigenfunctions [14] of the correlation matrix formed by independent samples of $\rho_{\mathrm{data}}(y)$. These eigenfunctions are automatically concentrated on the attractor.

We have implemented the program outlined here using "data" generated by the Hénon map of the plane (x_1, x_2) to itself,

$$x_1(n+1) = 1.0 - a x_1(n)^2 + x_2(n)\,,$$
$$x_2(n+1) = b x_1(n)\,, \tag{20}$$

with the familiar parameter values $a=1.4$ and $b=0.3$. Starting with data on $x_1(n)$, $n=1, 2, \ldots, N+1$ we constructed the two vectors $y(n)=(x_1(n), x_1(n+1))$ for $n=1, 2, \ldots, N$ as our basic data. In units set by the map itself, we found the minimum distance between points on the attractor to be $\approx 10^{-4}$ to 10^{-5} when $N \geq 500$. We chose the parameter σ in our map $F(y, a)$ to be about 100 times the square of this minimum. For large N this means that most points in the data set will have some neighbors. Specifically we took $\sigma = 3.4\times 10^{-7}$.

For our predictor we took three terms:

$$y(l+1) = X_1 F(y(l), a) + X_2 F^2(y(l-1), a) + X_3 F^3(y(l-2), a)\,. \tag{21}$$

Recall that we chose to fix the X_l rather than vary them in this note. To reflect our greater confidence in the lower-order iterates of $F(y, a)$, we took $X_1=0.8$ and $X_2=X_3=0.1$. Similar values of the X_l give much the same qualitative results. In the work we report here we have chosen the number of parameters a to be 4. We constrained the minimization of $C(X, a)$ over the a with λ_1 and with the mean of the phase function $\langle g\rangle$ so when both constraints are imposed we have two free variables among the a. The powers m_3 and m_4 in the $F(y, a)$ were chosen to be fixed during the minimization of $C(X, a)$. We took them to be $m_3=4$ and $m_4=5$. A parameter search over the m_k could have been done as well; but we have not done this.

There is an important item which we were required to address in our work. It is quite general, so we discuss it and our solution to it before reporting on particular calculations. The data we are given lies on the attractor; indeed, it defines the attractor. The attractor is an object of zero volume in R^D since its dimension is less than D. We have, thus, no information from the data on the behavior of the map $y \to F(y, a)$ when y lies in most of R^D. Our map must contain some rule which takes points y that are off the attractor and brings them onto the attracting set. Our class of maps $F(y, a)$ does that by mapping points that are off the attractor, which clearly have no neighbors among the data $y(n)$, to the origin $y=0$. We have addressed this matter by translating the origin of the coordinate system in which the $y(n)$ are given to lie well within $\sqrt{\sigma}$ of some data point. (Which data point we chose seemed not to matter.) When the parameters a have reached values near the optimum and that optimum is doing a good job of tracking the data, this translation of the origin is doing nothing. While we are searching the a, however, and are far from the optimum, this translation reinjects points mapped off the attractor by a bad $F(y, a)$ back onto or very near the attractor. This device, or an equivalent one, provides both stability and logic to the parameter search and is certainly needed for any $D \geq 2$.

If there were a way to probe the system producing

the data, we could avoid this need for reinjection by pulsing the probe and letting the system itself explore phase space off the attractor. The information we need would then be in the data set. The usual situation we anticipate is that the scalar time series $x(n)$ represents long-time behavior of the system and that only motion on the attractor is represented in the data.

In our work with data from the Hénon map we began with $N=750$. This was large enough to cover the attractor reasonably densely and allowed us to make calculations quickly. We also tried to choose N large enough to give an accurate picture of the attractor but small enough that we could imagine ourselves operating on a small data set we had been given. Our choice of N was partly motivated by our observations that the method we use to compute λ_1 when applied to the Hénon map directly undergoes some fluctuations for N much smaller than 400 and has settled down to a value $\lambda_1=0.408...$ after that. This value is consistent with other determinations of this Lyapunov exponent, so we were confident we had not chosen too small a value of N. With this value of N, we searched the parameters $\boldsymbol{a}$ to minimize $C(X, \boldsymbol{a})$. Our search utilized the software package NPSOL [15], which does not search for global minima, so some variation of initial values of the $\boldsymbol{a}$ was needed on our part. After some looking around we found a very shallow minimum in the variable a_2 near which the other a_i took the same values for large excursions in a_2. The values of the parameters at the minimum were $a_1=1.18$ $a_2=6.07$, $a_3=-0.0784$, and $a_4=0.0126$. At these values, the cost function $C(X, \boldsymbol{a})$ took the value 0.03967, while the Lyapunov exponent λ_1 was 5.95, rather than 0.408, and the value of the phase space average of $g(y)$ was 6.18 rather than the value of 2.81 computed from the data. Clearly we achieved a very good "fit" to the data as far as the cost function $C(X, \boldsymbol{a})$ was concerned, but the map $F(y, \boldsymbol{a})$ at that set of parameters $\boldsymbol{a}$ had little to do with the dynamical system generating the data.

Next we imposed only the Lyapunov constraint ($\lambda^{\text{map}}=\lambda^{\text{data}}$) on the minimization of the cost function. The parameter values found in our search were $a_1=1.19$, $a_2=0.816$, $a_3=0.000764$, and $a_4=-0.0121$. At this point the value of $C(X, \boldsymbol{a})$ was 0.04175, which is still an excellent "fit" in a least-squares sense. The Lyapunov exponent now differed from its value taken from the data by 0.05! This actually was a general rule we saw in our fitting of $F(y, \boldsymbol{a})$; namely, the smallest $C(X, \boldsymbol{a})$ did not give good values for the constraints, and a larger, though often not much larger, cost function was found by the constrained optimum. This means that while the constrained optimum will give a slightly worse point to point prediction of future values of points in R^D, it will, by construction, give better global properties.

Finally we imposed the equality of both the largest Lyapunov exponent ($\lambda^{\text{map}}=\lambda^{\text{data}}$) and the phase space average ($\langle g\rangle_{\text{map}}=\langle g\rangle_{\text{data}}$). The precise minimum resulting from the application of the search routine NPSOL depends on the acceptable limits we put on the satisfaction of the constraints. All minima were in the neighborhood of $C(X, \boldsymbol{a})=0.09$. For example, requiring each of the constraints to be met to an accuracy of 0.005, led to a cost function of 0.09082. Relaxing this to an accuracy of 0.01, led to a cost function of 0.09063, which is essentially the same. The parameter values shifted around in a common neighborhood for all these limits imposed on the constraints, indicating we were just moving around the same constrained minimum at various small distances. For constraints required to be satisfied within ± 0.006, the values of the a_i were $a_1=0.940$, $a_2=1.21$, $a_3=-0.0260$, and $a_4=0.00188$.

What is important here is not the specific set of values of the parameters $\boldsymbol{a}$. Rather, it is that we are able to find a_i's that meet both of our constraints *and* that although the a_i's that give accurate values of the constraints are quite different from those of the unconstrained minimum, they still give an acceptably small cost function.

We repeated this kind of calculation with a variety of values of N and on data sets generated with different initial conditions for the underlying Hénon map. For $N=1200$ and $N=1700$, for example, we report in table 1 the results of calculations precisely along the lines just discussed. They are rather similar in character to the results for $N=750$ just reported and to the results for other values of N we explored. In the table we show the $C(X, \boldsymbol{a})$ for both the constrained and unconstrained optimization. The values of the constraints λ_1 and $\langle g\rangle$ are shown for the map both when those constraints have been imposed and when they have not.

Table 1
Optimization results. $C(X, a)$ (eq. (9) with $F(y, a)$ from eq. (4)) is shown with and without invariant constraints. The following parameters are used for all N: $X_1=0.8$, $X_2=0.1$, $X_3=0.1$, $\lambda_1^{data}=0.408$.

N	$\langle g\rangle^{data}$		$C(X,a)$	a_1	a_2	a_3	a_4	λ_1^{map}	$\langle g\rangle^{map}$
750	2.81	unconst.	0.03502	1.18	607	−0.0784	0.0126	5.95	6.18
		λ_1	0.04175	1.19	0.816	0.000764	−0.0121	0.403	6.16
		$\lambda_1, \langle g\rangle$	0.09105	0.940	1.21	−0.0260	0.00188	0.406	2.82
1200	2.81	unconst.	0.03090	1.19	14.30	−0.0650	0.00912	2.72	6.24
		λ_1	0.03113	1.17	0.677	−0.0645	0.00961	0.408	5.96
		$\lambda_1, \langle g\rangle$	0.0851	0.966	−0.718	−0.00757	−0.00649	0.411	2.81
1700	2.90	unconst.	0.03444	1.16	−13.36	−0.0809	0.00150	2.50	6.48
		λ_1	0.03894	1.17	−0.246	−0.126	0.00301	0.405	6.32
		$\lambda_1, \langle g\rangle$	0.0926	0.939	0.478	−0.0355	−0.00237	0.409	2.91

We have now demonstrated that using a predictor of the form

$$y(N+1)=\sum_{j=1}^{L} X_j F^j(y(N-j+1, a)$$

with our class of mapping functions $F(y, a)$ can give excellent least-squares fits to chaotic time series data on a strange attractor while simultaneously satisfying constraints on those fits dictated by the geometrical invariants which characterize that attractor. Furthermore, a straightforward least-squares fit to the data does not, in general, reproduce the dynamical information on Lyapunov exponents and invariant densities that are contained in the data itself. To produce the correct value of the invariants we must accept a larger cost function. This loss in least-squares based predicted power is made up for by the built in quality of our predictors; namely, they will produce the correct long-term *statistical* behavior of the dynamical system whose properties we are trying to learn by the analysis of the original scalar time series.

Several directions are clear for further investigation. One is to apply this set of methods to higher-dimensional systems both for computer generated data and for experimental data. Before the latter is addressed we must study in detail the important issue of how to treat extrinsic noise within the kind of phase space description of time series we are using here. Some ideas on that are contained in other work [9,16], though no consideration of it has been given here. Another quite interesting issue is the development of a set of reliable and efficient algorithms for the extraction of invariant quantities such as Lyapunov exponents from data. Finally the extension of the methods demonstrated here to systems with spatial degrees of freedom would be most interesting. We plan to discuss many of these items in our own expansion of this short note [7].

We are most appreciative for productive conversations with K. Bruckner, M. Freedman, H. Levine, J. Theiler, and Bruce West about the material covered in this note. This work was supported in part under a contract with the DARPA Applied and Computational Mathematics Program, No. F 49620-87-C-0117 and in part under the DARPA-University Research Initiative, URI Contract No. N00014-86-K-0758. J.B. Kadtke wishes to acknowledge support of AFOSR grant No. AFOSR-89-0072.

References

[1] J.P. Eckmann and D. Ruelle, Rev. Mod. Phys. 57 (1985) 617.
[2] N.H. Packard, J.P. Crutchfield, J.D. Farmer and R.S. Shaw, Phys. Rev. Lett. 45 (1980) 712.
[3] F. Takens, in: Lecture notes in mathematics, Vol. 898. Dynamical systems and turbulence, Warwick 1980, eds. D. Rand and L.S. Young (Springer, Berlin, 1981) pp. 366–381.
[4] R. Mañé, in: Lecture notes in mathematics, Vol. 898. Dynamical systems and turbulence, Warwick 1980, eds. D. Rand and L.S. Young, (Springer, Berlin, 1981) pp. 230–242.

[5] A.M. Fraser and H.L. Swinney, Phys. Rev. A 33 (1986) 1134; Information and entropy in strange attractors, PhD Dissertation, University of Texas at Austin (May 1988).

[6] L.-S. Young, Ergod. Theory Dynam. Syst. 2 (1982) 109; P. Grassberger and I. Procaccia, Phys. Rev. Lett. 50 (1983) 346; Physica D 9 (1983) 189; J. Theiler, Quantifying chaos: practical estimation of the correlation dimension, PhD Dissertation, California Institute of Technology (June 1987); preprints from the UCSD Institute for Nonlinear Science (June 1988).

[7] H.D.I. Abarbanel, R. Brown and J.B. Kadtke, Prediction, system identification, and control on strange attractors: methods for time series with broadband power spectra, UCSD Institute for Nonlinear Science Preprint (January 1989).

[8] M. Hénon, Commun. Math. Phys. 50 (1976) 69; E.N. Lorenz, J. Atmos. Sci. 20 (1963) 130.

[9] J.D. Farmer and J.J. Sidorowich, Phys. Rev. Lett. 59 (1987) 845; Exploiting chaos to predict the future and reduce noise, Preprint LA-UR-88-901, Center for Nonlinear Studies, Los Alamos National Laboratory (March 1988); A.S. Lapedes and R. Farber, Nonlinear signal processing using neural networks: prediction and system modeling, Preprint LA-UR-87-2662, Los Alamos National Laboratory (1987); How neural nets work, Preprint LA-UR-88-418, Los Alamos National Laboratory (January 1988); J.P. Crutchfield and B.S. McNamara, Complex Syst. 1 (1987) 417.

[10] I.J. Leontaritis and S.A. Billings, Int. J. Control 41 (1985) 303.

[11] W.H. Press, B.P. Flannery, S.A. Teukolsky and W.T. Vetterling, Numerical recipes (Cambridge Univ. Press, Cambridge, 1986) section 12.10.

[12] V.I. Oseledec, Tr. Mosk. Mat. Obsc. Moscow Math. Soc. 19 (1968) 17.

[13] I. Shimada and T. Nagashima, Prog. Theor. Phys. 61 (1979) 1605; G. Benettin, C. Froeschle and J.P. Scheidecker, Phys. Rev. A 19 (1979) 2454; A. Wolf, J.B. Swift, H.L. Swinney and J.A. Vastano, Physica D 16 (1985) 285; J.-P. Eckmann, S.O. Kamphorst, D. Ruelle and S. Ciliberto, Phys. Rev. A 34 (1986) 4971.

[14] A. Rosenfeld and A.C. Kak, Digital picture processing, 2nd Ed. (Academic Press, New York, 1981) chs. 4, 5.

[15] Ph.E. Gill, W. Murray, M.A. Saunders and M.H. Wright, User's Guide for NPSOL (Version 4.0): a FORTRAN package for nonlinear programming, Technical Report SOL 86-2 (January 1986), Systems Optimization Laboratory, Stanford University.

[16] Yu.A. Kravtsov and V.S. Etkin, Izv. Vyssh. Uchebn. Zaved. Radiofiz. 24 (1981) 992 [Radiophys. Quantum Electron. 24 (1982) 679].

Prediction in chaotic nonlinear systems: Methods for time series with broadband Fourier spectra

Henry D. I. Abarbanel
Institute for Nonlinear Science, University of California, San Diego, La Jolla, California 92093-0402; Marine Physical Laboratory, Scripps Institution of Oceanography, La Jolla, California 92093-0701; and Department of Physics, University of California, San Diego, La Jolla, California 92093-0319

Reggie Brown and James B. Kadtke
Institute for Nonlinear Science, University of California, San Diego, La Jolla, California 92093-0402
(Received 22 May 1989; revised manuscript received 9 August 1989)

We consider the problem of prediction and system identification for time series having broadband power spectra that arise from the intrinsic nonlinear dynamics of the system. We view the motion of the system in a reconstructed phase space that captures the attractor (usually strange) on which the system evolves and give a procedure for constructing parametrized maps that evolve points in the phase space into the future. The predictor of future points in the phase space is a combination of operation on past points by the map and its iterates. Thus the map is regarded as a dynamical system and not just a fit to the data. The invariants of the dynamical system, the Lyapunov exponents and optimum moments of the invariant density on the attractor, are used as constraints on the choice of mapping parameters. The parameter values are chosen through a constrained least-squares optimization procedure, constrained by the values of these invariants. We give a detailed discussion of methods to extract the Lyapunov exponents and optimum moments from data and show how to equate them to the values for the parametric map in the constrained optimization. We also discuss the motivation and methods we utilize for choosing the form of our parametric maps. Their form has a strong similarity to the work in statistics on kernel density estimation, but the goals and techniques differ in detail. Our methodology is applied to "data" from the Hénon map and the Lorenz system of differential equations and shown to be feasible. We find that the parameter values that minimize the least-squares criterion do not, in general, reproduce the invariants of the dynamical system. The maps that do reproduce the values of the invariants are not optimum in the least-squares sense, yet still are excellent predictors. We discuss several technical and general problems associated with prediction and system identification on strange attractors. In particular, we consider the matter of the evolution of points that are off the attractor (where few or no data are available), onto the attractor where long-term motion takes place. We find that we are able to realize maps that give a least-squares approximation to the data with rms variation over the attractor of 0.5% or less and still reproduce the dynamical invariants to 5% or better. The dynamical invariants are the classifiers of the dynamical system producing the broadband time series in the first place, so this quality of the maps is essential in representing the correct dynamics.

I. INTRODUCTION

A. General remarks

Analysis of time series from dynamical systems is an important issue in many different fields of engineering and science. The most common tool for this analysis is the Fourier (or other similar) transform of the data $x(n)$ to discover sharp lines in its power spectrum. Spectral identification lies at the heart of much of the work on *linear* systems to which time series analysis is applied.[1,2] When one encounters a broadband power spectrum, the common assumption is that it represents *extrinsic noise* and not characteristics of the signal.

It has become increasingly clear in recent years that nonlinear systems exhibiting deterministic chaos will generate a time series whose power spectrum is broadband. Generically, dissipative nonlinear chaotic systems evolve nonperiodically on a strange attractor that lives in a phase space of finite (and often small) dimension. *Noise* does not evolve on a strange attractor and will occupy an arbitrarily large number of dimensions. Hence to model nonlinear chaotic systems as noise is certainly incorrect. For these systems the source of the broadband spectrum is the *intrinsic* chaotic dynamics that underlies the time series.

Our focus in this work is on signals with a substantial broadband power spectrum which, since external noise is absent or very small, represents the nonperiodic behavior of a dynamical system whose orbits lie on a strange attractor. The idea, now rather well established, that such an object can have a small fractal dimension (and still govern the long time evolution of a system with far more numerous degrees of freedom than represented by the dimension of the attractor) is really the starting point of our work.[3,4]

 ©1990 The American Physical Society

It is very important that though $x(n)$ may be a long, quiet data set it is likely to have a very broad power spectrum. Indeed, if the signal one is studying has a power spectrum with substantial strong lines, one is well advised to recognize the implied sinusoids as the underlying linear degrees of freedom and avoid altogether the labor we propose here.

It has been shown that in nonlinear systems that exhibit deterministic chaos one can determine from the observation of a *single* dynamical variable the geometric structure of the *many* variable dynamics that produced the measured signal.[5–10] The method that has developed for the construction of the phase space in which the dynamics dwells is called *phase-space reconstruction*. The result of this reconstruction is an embedding space of d dimensions (d is an integer) in which one may observe the attractor. One can view the evolution in the reconstructed phase space of the many dimensional dynamics in a quantitative fashion in the time domain.

In this article we describe both in outline and implementation a program for extracting from the observations of this single broadband temporal signal quantitative predictions for the evolution of initial conditions differing from the observed data points. We assume that once transients are gone the evolution of the system is on a strange attractor with dimension d_A, where d_A is generally fractional. If the evolution of the system is on such an attractor, then the d-dimensional embedding space enclosing the attractor should be sufficiently larger than d_A that all the geometric information about the attractor is exposed in the embedding space. Mañé and Takens's[6,7] formal result requires $d > 2d_A + 1$ to assure one of a faithful representation of the d_A-dimensional attractor as seen in the d-dimensional embedding space, but often, in practice, $d > d_A$ will do. The method of phase-space reconstruction seeks to construct from the $x(n)$'s d-dimensional vectors which, when embedded in $\mathbf{R}^d$ describes the full dynamical evolution of the system. Section II is devoted to the issue of identifying the correct value of d from the data set.

For the moment suppose we have found d by one means or another. We imagine measuring a single scalar variable x at discrete time points $x(n)$ for $n=1,2,\ldots,N_D$. (Observation of several dynamical variables from the system is even better, and serves to provide confirmation of the information on the deductions from observations of any single variable.) We can construct d-dimensional vectors $\mathbf{y}(n)$ in the embedding space by

$$\mathbf{y}(n)=(x(n),x(n+\tau_1),x(n+\tau_2),\ldots,x(n+\tau_{d-1}))\ ,$$

for some set of time lags $\tau_1,\tau_2,\ldots,\tau_{d-1}$. The set of $\mathbf{y}(n)$'s, of which we have $N=N_D-d$, capture the evolution of the nonlinear system under observation as it moves through the d-dimensional phase space. Familiar phase-space coordinates are the time derivatives $x(n),dx(n)/dt,d^2x(n)/dt^2,\ldots$, evaluated at discrete times. The data on $x(n)$ are acquired only at discrete times and establishing the values of these derivatives is certain to be inaccurate. The time lagged $x(n)$'s, which are the coordinate elements of the $\mathbf{y}(n)$'s, are nonlinear combinations of the local time derivatives and are fully acceptable substitutes for the usual phase-space coordinates. This has been emphasized by Eckmann and Ruelle.[10]

With the $\mathbf{y}(n)$'s and the embedding space in hand, we ask here the ambitious question of how we can use the series of $\mathbf{y}(n)$'s to predict $\mathbf{y}(N+1),\mathbf{y}(N+2)$, etc. Equivalently, we can ask what is the evolution, under the same dynamical system that produced the $\mathbf{y}(n)$'s, of a point $\mathbf{y}$, that is on the attractor but not in the original data set. We will have answered this question when given a data set $\mathbf{y}(1),\mathbf{y}(2),\ldots,\mathbf{y}(N)$, we have identified a "reliable" map $\mathbf{F}$ from $\mathbf{R}^d$ to itself parametrized by $\mathbf{a}=(a_1,a_2,\ldots,a_p)$ which takes us from $\mathbf{y}(n)$ to $\mathbf{y}(n+1)$,

$$\mathbf{y}(n+1)=\mathbf{F}(\mathbf{y}(n),\mathbf{a})\ .$$

If we can establish a reliable $\mathbf{F}(\mathbf{y},\mathbf{a})$, then the evolution of a point $\mathbf{y}$ in $\mathbf{R}^d$ that is not a member of the measured data set would be $\mathbf{y}\rightarrow\mathbf{y}_1=\mathbf{F}(\mathbf{y},\mathbf{a})$, $\mathbf{y}_1\rightarrow\mathbf{y}_2=\mathbf{F}(\mathbf{y}_1,\mathbf{a})=\mathbf{F}(\mathbf{F}(\mathbf{y},\mathbf{a}),\mathbf{a})=\mathbf{F}^2(\mathbf{y},\mathbf{a})$, etc.

Our first view of the data $\mathbf{y}(1),\mathbf{y}(2),\ldots,\mathbf{y}(N)$ is that it can be thought of as a pair of columns of vectors in $\mathbf{R}^d$

$$\begin{array}{ll}\mathbf{y}(1) & \mathbf{y}(2)\\ \mathbf{y}(2) & \mathbf{y}(3)\\ \vdots & \vdots\\ \mathbf{y}(n) & \mathbf{y}(n+1)\\ \vdots & \vdots\\ \mathbf{y}(N-1) & \mathbf{y}(N)\ ,\end{array}$$

and our function $\mathbf{F}(\mathbf{y},\mathbf{a})$ comes from parametrically "fitting" the right-hand column of $\mathbf{y}(n+1)$ resulting from the left-hand column of $\mathbf{y}(n)$. Fitting the data then suggests making a least-squares estimation of $\mathbf{a}$ so that the *cost function*

$$\hat{C}(\mathbf{a})=\sum_{n=1}^{N-1}\left\{\sum_{m=1}^{d}[y_m(n+1)-F_m(\mathbf{y}(n),\mathbf{a})]^2\right\}$$

is minimized. Our approach differs from previous work in detailed tactics and in our imposition of important geometrical structure as constraints on the minimization of the cost function. The articles we have greatly relied on for guidance and initial impetus in our research are those by Farmer and Sidorovitch[11] (we refer to this paper as FS in the following), Lapedes and Farber,[12] and Crutchfield and McNamara.[13]

Our main point, simply stated, is that we are not just making a fit to data with a set of functions $\mathbf{F}(\mathbf{y},\mathbf{a})$. Rather, these functions evaluated along the orbit are to be related to each other in the manner of a dynamical system. This leads to a rather different view of the fitting functions than the one usually taken in trying to match data to observations. It means that the function $\mathbf{F}(\mathbf{y},\mathbf{a})$ evalu-

It is very important that though $x(n)$ may be a long, quiet data set it is likely to have a very broad power spectrum. Indeed, if the signal one is studying has a power spectrum with substantial strong lines, one is well advised to recognize the implied sinusoids as the underlying linear degrees of freedom and avoid altogether the labor we propose here.

It has been shown that in nonlinear systems that exhibit deterministic chaos one can determine from the observation of a *single* dynamical variable the geometric structure of the *many* variable dynamics that produced the measured signal.[5–10] The method that has developed for the construction of the phase space in which the dynamics dwells is called *phase-space reconstruction*. The result of this reconstruction is an embedding space of d dimensions (d is an integer) in which one may observe the attractor. One can view the evolution in the reconstructed phase space of the many dimensional dynamics in a quantitative fashion in the time domain.

In this article we describe both in outline and implementation a program for extracting from the observations of this single broadband temporal signal quantitative predictions for the evolution of initial conditions differing from the observed data points. We assume that once transients are gone the evolution of the system is on a strange attractor with dimension d_A, where d_A is generally fractional. If the evolution of the system is on such an attractor, then the d-dimensional embedding space enclosing the attractor should be sufficiently larger than d_A that all the geometric information about the attractor is exposed in the embedding space. Mañé and Takens's[6,7] formal result requires $d > 2d_A + 1$ to assure one of a faithful representation of the d_A-dimensional attractor as seen in the d-dimensional embedding space, but often, in practice, $d > d_A$ will do. The method of phase-space reconstruction seeks to construct from the $x(n)$'s d-dimensional vectors which, when embedded in $\mathbf{R}^d$ describes the full dynamical evolution of the system. Section II is devoted to the issue of identifying the correct value of d from the data set.

For the moment suppose we have found d by one means or another. We imagine measuring a single scalar variable x at discrete time points $x(n)$ for $n=1,2,\ldots,N_D$. (Observation of several dynamical variables from the system is even better, and serves to provide confirmation of the information on the deductions from observations of any single variable.) We can construct d-dimensional vectors $\mathbf{y}(n)$ in the embedding space by

$$\mathbf{y}(n)=(x(n),x(n+\tau_1),x(n+\tau_2),\ldots,x(n+\tau_{d-1}))\ ,$$

for some set of time lags $\tau_1,\tau_2,\ldots,\tau_{d-1}$. The set of $\mathbf{y}(n)$'s, of which we have $N=N_D-d$, capture the evolution of the nonlinear system under observation as it moves through the d-dimensional phase space. Familiar phase-space coordinates are the time derivatives $x(n),dx(n)/dt,d^2x(n)/dt^2,\ldots$, evaluated at discrete times. The data on $x(n)$ are acquired only at discrete times and establishing the values of these derivatives is certain to be inaccurate. The time lagged $x(n)$'s, which are the coordinate elements of the $\mathbf{y}(n)$'s, are nonlinear combinations of the local time derivatives and are fully acceptable substitutes for the usual phase-space coordinates. This has been emphasized by Eckmann and Ruelle.[10]

With the $\mathbf{y}(n)$'s and the embedding space in hand, we ask here the ambitious question of how we can use the series of $\mathbf{y}(n)$'s to predict $\mathbf{y}(N+1),\mathbf{y}(N+2)$, etc. Equivalently, we can ask what is the evolution, under the same dynamical system that produced the $\mathbf{y}(n)$'s, of a point $\mathbf{y}$, that is on the attractor but not in the original data set. We will have answered this question when given a data set $\mathbf{y}(1),\mathbf{y}(2),\ldots,\mathbf{y}(N)$, we have identified a "reliable" map $\mathbf{F}$ from $\mathbf{R}^d$ to itself parametrized by $\mathbf{a}=(a_1,a_2,\ldots,a_P)$ which takes us from $\mathbf{y}(n)$ to $\mathbf{y}(n+1)$,

$$\mathbf{y}(n+1)=\mathbf{F}(\mathbf{y}(n),\mathbf{a})\ .$$

If we can establish a reliable $\mathbf{F}(\mathbf{y},\mathbf{a})$, then the evolution of a point $\mathbf{y}$ in $\mathbf{R}^d$ that is not a member of the measured data set would be $\mathbf{y}\rightarrow\mathbf{y}_1=\mathbf{F}(\mathbf{y},\mathbf{a})$, $\mathbf{y}_1\rightarrow\mathbf{y}_2=\mathbf{F}(\mathbf{y}_1,\mathbf{a})=\mathbf{F}(\mathbf{F}(\mathbf{y},\mathbf{a}),\mathbf{a})=\mathbf{F}^2(\mathbf{y},\mathbf{a})$, etc.

Our first view of the data $\mathbf{y}(1),\mathbf{y}(2),\ldots,\mathbf{y}(N)$ is that it can be thought of as a pair of columns of vectors in $\mathbf{R}^d$

$$\begin{array}{ll}\mathbf{y}(1) & \mathbf{y}(2)\\ \mathbf{y}(2) & \mathbf{y}(3)\\ \vdots & \vdots\\ \mathbf{y}(n) & \mathbf{y}(n+1)\\ \vdots & \vdots\\ \mathbf{y}(N-1) & \mathbf{y}(N)\ ,\end{array}$$

and our function $\mathbf{F}(\mathbf{y},\mathbf{a})$ comes from parametrically "fitting" the right-hand column of $\mathbf{y}(n+1)$ resulting from the left-hand column of $\mathbf{y}(n)$. Fitting the data then suggests making a least-squares estimation of $\mathbf{a}$ so that the *cost function*

$$\hat{C}(\mathbf{a})=\sum_{n=1}^{N-1}\left[\sum_{m=1}^{d}[y_m(n+1)-F_m(\mathbf{y}(n),\mathbf{a})]^2\right]$$

is minimized. Our approach differs from previous work in detailed tactics and in our imposition of important geometrical structure as constraints on the minimization of the cost function. The articles we have greatly relied on for guidance and initial impetus in our research are those by Farmer and Sidorovitch[11] (we refer to this paper as FS in the following), Lapedes and Farber,[12] and Crutchfield and McNamara.[13]

Our main point, simply stated, is that we are not just making a fit to data with a set of functions $\mathbf{F}(\mathbf{y},\mathbf{a})$. Rather, these functions evaluated along the orbit are to be related to each other in the manner of a dynamical system. This leads to a rather different view of the fitting functions than the one usually taken in trying to match data to observations. It means that the function $\mathbf{F}(\mathbf{y},\mathbf{a})$ evalu-

ated on the data vector $\mathbf{y}(n)$ is required to do more than reproduce $\mathbf{y}(n+1)$ as accurately as possible. $\mathbf{F}(\mathbf{y},\mathbf{a})$ must also be a function which when iterated will reproduce $\mathbf{y}(n+2)$ after two applications to $\mathbf{y}(n)$ and $\mathbf{y}(n+3)$ after three, etc. The notion of $\mathbf{F}(\mathbf{y},\mathbf{a})$ as a dynamical system also leads to modifications of the cost function. The cost function should reflect the fact that iterations of $\mathbf{F}(\mathbf{y},\mathbf{a})$ also yield points on the orbit. Furthermore, under our approach geometrical properties of the dynamical system given by $\mathbf{F}(\mathbf{y},\mathbf{a})$ are used to determine the success of the fit. It is not just the function's ability to reproduce in a least-squares sense the observed data that is important. The data contain invariant information that is essential for a full description of the geometrical structure of the attractor that it evolves on. Our key observation in this article is that, in general, least-squares fitting alone does not produce a map that captures the invariant characteristics of the attractor described by the data $\mathbf{y}(n)$, $n=1,\ldots,N$. One must calculate from the data as many of these invariant quantities as possible and then impose them as constraints on the fit. In this way we emphasize the fact that one is creating a dynamics and not just a fit to data. The product of our minimization of the constrained cost function is a mapping $\mathbf{F}(\mathbf{y},\mathbf{a})$ of $\mathbf{R}^d$ to itself which is not only *reliable* in that it reproduces the given data set by having a small cost function, but is also *representational* in that it has the same geometric invariants as the underlying dynamical system. The methods for identifying those invariants and utilizing them as classifiers for the dynamical system is a matter of some importance in itself.

The invariants are properties of the function $\mathbf{F}(\mathbf{y},\mathbf{a})$ viewed as a dynamical system which maps $\mathbf{R}^d$ to itself. We will concentrate on two kinds of invariants. One kind of invariant, the Lyapunov characteristic exponents $\lambda_1,\lambda_2,\ldots,\lambda_d$, describes the expansion or contraction of phase-space volumes under the iteration of $\mathbf{F}(\mathbf{y},\mathbf{a})$.[10–17] Lyapunov exponents are invariant under smooth changes of coordinate and are independent of the initial conditions of the orbit one follows on the attractor. The second kind of invariant is the density of points on the attractor $\rho(\mathbf{y})$. It captures *global* features of the frequency with which orbits visit various portions of the attractor. The density is a different kind of invariant than the Lyapunov exponents. Its integrals with smooth functions $G(\mathbf{y})$ are unchanged under operation with the mapping function which underlies the dynamics $\mathbf{y}(n)\rightarrow\mathbf{y}(n+1)$,

$$\int d^d y\, \rho(\mathbf{y})G(\mathbf{y})=\int d^d y\, \rho(\mathbf{y})G(\mathbf{F}(\mathbf{y},\mathbf{a}))\,.$$

It too is independent of the initial conditions on the orbits.[10,18,19]

In this paper we find the parameters $\mathbf{a}$ in $\mathbf{F}(\mathbf{y},\mathbf{a})$ by minimizing a cost function subject to certain constraints. The constraints are chosen to insure that iterations of the mapping function $\mathbf{F}(\mathbf{y},\mathbf{a})$ give rise to values of dynamical invariants which are the same as those indicated by the experimentally measured data set $\mathbf{y}(n)$. In this way essential geometric information about the particular attractor on which the data live will be built into the parametric mapping. Straightforward least-squares minimization does not accurately reproduce these invariants. Thus one must perform a least-square minimization subject to the constraints that $\mathbf{F}(\mathbf{y},\mathbf{a})$ accurately produce the Lyapunov spectra $\lambda_1,\lambda_2,\ldots,\lambda_d$ and the invariant density $\rho(\mathbf{y})$. This paper is devoted to explaining in detail how one implements the idea just stated.

B. Choosing maps and predictors

Assuming for the moment that we have successfully embedded the data $x(n)$ in $\mathbf{R}^d$ by creating d-dimensional vectors $\mathbf{y}(n)$, $n=1,\ldots,N$. We need to choose a class of parametrized mappings, a cost function to minimize, and a means to impose the constraints on our minimization. The maps must have some way of fitting the data by closely reproducing one data point from the previous one by $\mathbf{y}(n+1)\approx\mathbf{F}(\mathbf{y}(n),\mathbf{a})$. Our maps are required to "look around" at the behavior of the phase-space neighbors of the point $\mathbf{y}(n)$ and predict forward according to how a cluster of phase-space neighbors, regardless of their temporal sequence, are moved forward in time. The idea here is that one may use knowledge of the behavior of local regions of phase space as well as past points on an orbit to determine where a point will be mapped in the temporal future. The maps we choose must then be sensitive to their neighborhood in phase space and must inquire about the fate of any spatial neighbor under the map without concern of its temporal arrival in the neighborhood. The map will then try to take any new point $\mathbf{y}$ and map it forward to some weighted average of its neighbors' forward evolution.

We take our mappings to be of the form

$$\mathbf{F}(\mathbf{y},\mathbf{a})=\sum_{n=1}^{N-1}\mathbf{y}(n+1)g(\mathbf{y},\mathbf{y}(n);\mathbf{a})\,, \tag{1}$$

where $g(\mathbf{y},\mathbf{y}(n);\mathbf{a})$ is near 1 for $\mathbf{y}=\mathbf{y}(n)$, and vanishes rapidly for nonzero $|\mathbf{y}-\mathbf{y}(n)|$; the vertical bars represent some norm, in our case Euclidean, in $\mathbf{R}^d$. $\mathbf{F}(\mathbf{y}(k),\mathbf{a})$ will then be quite close to $\mathbf{y}(k+1)$.

This type of mapping is strongly suggestive of the form used in the statistical literature under the name of *kernel estimation* or *kernel density estimation*. An explicit recent example that illustrates the similarity is found in Ref. 20. Other useful discussions of this method applied to various problems are to be found in Refs. 21 and 22; our attention was directed to this similarity by Farmer and Sidorowich.[23] We do not claim to have a better method for choosing our function g than those in the literature, but our motivation here does differ from all the citations except Rice.[20] Our constraints on g are also different, but could be modified. For example, the integral of g over $\mathbf{y}$ need not be unity, nor do we require that g be positive. We will return to a discussion of choices for g in our summary in Sec. VI.

Our choice here for $g(\mathbf{y},\mathbf{y}(n);\mathbf{a})$—one among many, of course—is this:

$$g(\mathbf{y},\mathbf{y}(n);\mathbf{a})=\frac{\exp[-|\mathbf{y}-\mathbf{y}(n)|^2/\sigma]\left[a_1+a_2\mathbf{y}(n)\cdot(\mathbf{y}-\mathbf{y}(n))+\sum_{k=3}^{P}a_k(|\mathbf{y}-\mathbf{y}(n)|^2/\sigma)^{m_k}\right]}{\sum_{n=1}^{N-1}\exp[-|\mathbf{y}-\mathbf{y}(n)|^2/\sigma]\left[a_1+\sum_{k=3}^{P}a_k(|\mathbf{y}-\mathbf{y}(n)|^2/\sigma)^{m_k}\right]} . \quad (2)$$

The parameter space $\mathbf{a}$ is P dimensional, $\mathbf{a}=(a_1,a_2,\ldots,a_P)$. σ is a fixed parameter that provides a scale we can use to determine which points in the data set are "close" to $\mathbf{y}$. The m_k's are also fixed at various values. We could treat both σ and the m_k's as parameters to be optimized in the same sense as the $\mathbf{a}$'s. However, we choose not to do this in our work; not for any fundamental reasons, but because we wished to explore other issues and wished to keep down the size of the parameter space over which our minimization searches were performed.

The weight function $g(\mathbf{y},\mathbf{y}(n);\mathbf{a})$ which we use was arrived at after some experimentation. It, as do many other choices, certainly satisfies our general requirements. These requirements include the following:

The function is sensitive to the presence of near "neighbors" in phase space. Only points $\mathbf{y}(n)$ within a distance from $\mathbf{y}$ of order $\sqrt{\sigma}$ make any sizable contribution to $g(\mathbf{y},\mathbf{y}(n);\mathbf{a})$.

When $\sigma\to 0$, $g(\mathbf{y},\mathbf{y}(n);\mathbf{a})$ becomes essentially a Kronecker delta and the point $\mathbf{y}(n)$ is mapped precisely to $\mathbf{y}(n+1)$.

It is easy to differentiate both in $\mathbf{y}$ and in $\mathbf{a}$. These derivatives are important in the minimization of the cost function using our methods, and having explicit expressions for the required derivatives in either of these independent variables makes the optimization routines run much faster.

In the function we have chosen it is easy to retain many parameters all of the same general form, thus as the number of constraints on the optimization of the cost function is increased, the pattern of our searches remains the same.

The essential function which senses neighbors, namely the exponential, can easily be replaced by other choices, such as those in Table 3.1 of Silverman's monograph.[21] The general form of our arguments goes through then without modification.

By virtue of the term involving a_2 in the numerator, this form of $g(\mathbf{y},\mathbf{y}(n);\mathbf{a})$ allowed us to satisfy constraints set by the Lyapunov exponents with numerical stability and accuracy. The denominator serves as an approximate counter for the number of neighbors of the point $\mathbf{y}$, so the numerator works less to produce the required average for the forward prediction of the point $\mathbf{y}$. The presence of the denominator assured us of numerical ease in making the parameters in the map $\mathbf{F}(\mathbf{y},\mathbf{a})$ meet our requirement of producing an average over neighborhood points in projecting forward in time any phase-space point. This made the numerical algorithms we use much more efficient and accurate.

The choice of cost function is also rather much up to us. Since we are to think of $\mathbf{F}(\mathbf{y},\mathbf{a})$ as a dynamical system evolving points $\mathbf{y}(n)$ into new points $\mathbf{y}(n+1)$, we should consider asking the map to reproduce accurately from $\mathbf{y}(n)$ not only the "next" point $\mathbf{y}(n+1)$ but, via iteration, a sequence of points $\mathbf{y}(n+1),\mathbf{y}(n+2),\mathbf{y}(n+3),\ldots,\mathbf{y}(n+L)$ up to some L beyond which we simply do not trust the accuracy of our algorithm $\mathbf{F}$ or of the machines we use to compute the future $\mathbf{y}$'s.

This suggests the *predictor* for future points to be a linear combination of iterated powers of the map $\mathbf{F}(\mathbf{y},\mathbf{a})$,

$$\mathbf{y}(m+1)=\sum_{k=1}^{L}X_k\mathbf{F}^k(\mathbf{y}(m-k+1),\mathbf{a}) , \quad (3)$$

where $\mathbf{F}^k$ is the kth iterate of $\mathbf{F}$ as described above. If $\mathbf{F}(\mathbf{y},\mathbf{a})$ were the exact mapping, then each term in the sum over k would be $X_k\mathbf{y}(m+1)$. Thus we require

$$\sum_{k=1}^{L}X_k=1 .$$

The X's weight the various iterates of $\mathbf{F}$ and are used to determine which iterates of $\mathbf{F}$ we believe are the most accurate. Typically, one would require $X_j\geq X_{j+1}$ to indicate that the lower iterates of $\mathbf{F}$ are believed to be more accurate than the higher iterates. This predictor is a natural generalization to the nonlinear problem of the common linear predictor

$$\mathbf{y}(m+1)=\sum_{k=1}^{L}X_k\mathbf{y}(m-k+1) ,$$

with the clear differences associated with the iterative nature of the map $\mathbf{F}(\mathbf{y},\mathbf{a})$.

This predictor [Eq. (3)] combines both past temporal information from times $m-k+1$; $k=1,2,\ldots,L$ and information from all the phase-space neighbors of the orbit points $\mathbf{y}(m-k+1)$ because of the structure of $\mathbf{F}(\mathbf{y},\mathbf{a})$. The combination of spatial and temporal information provides a significant "lever arm" which permits Eq. (3) to quite accurately make forecasts about the forward evolution of points $\mathbf{y}$ in $\mathbf{R}^d$. By utilizing the phase-space information in $\mathbf{F}(\mathbf{y},\mathbf{a})$ at each temporal step we efficiently tap properties of the full data set.

The cost function associated with this predictor is

$$C(\mathbf{X},\mathbf{a})=\frac{\sum_{n=L}^{N-1}\left[|\mathbf{y}(n+1)-\sum_{k=1}^{L}X_k\mathbf{F}^k(\mathbf{y}(n-k+1),\mathbf{a})|^2\right]}{\sum_{n=1}^{N}|\mathbf{y}(n)\cdot\mathbf{y}(n)|^2} . \quad (4)$$

This kind of cost function will automatically contain information on the Lyapunov exponents which themselves are expressions of the dynamics as iterations of the map.

Some information on the invariant density function on the attractor is also contained in this improved cost function.[24]

Another major consideration to us is the great difference in the coordinate scale of various attractors. The numerator of the cost function [Eq. (4)] is the residual of the mapped function summed over the entire trajectory, and hence gives a measure of the sum of the absolute errors over all the mapped points. Since the absolute error is obviously dependent on the macroscale of the attractor, it is more informative to rescale the final cost function value in some manner which reflects error. In our samples, the scale of the attractor of the Hénon attractor is on the order of unity, while that of the Lorenz attractor is of order 100. Hence some form of rescaling of the cost function became desirable in order to have a relative measure of comparison between two systems with different macroscale. We chose a normalization in the following straightforward manner: we simple summed the magnitudes of the position vectors of all the points on the attractor, and retained this value as a constant. Absolute values for the cost function after normalization by the denominator in Eq. (4) give a more sensible relative measure of the error of our prediction function $\mathbf{F}(\mathbf{y},\mathbf{a})$.

We note that FS suggest forecasting the evolution of a point $\mathbf{y}$ by looking around at the neighbors of $\mathbf{y}$ among the data set $\mathbf{y}(n)$ and observing where these neighbors go under one iteration of the underlying map taking the $\mathbf{y}(n)$ to $\mathbf{y}(n+1)$. They determine the future of the new point $\mathbf{y}$ by an interpolation involving the future of its neighbors. Our mapping function Eqs. (1) and (2) does precisely this as indicated. All points in the data set are given some weight in the future of $\mathbf{y}$, but if $g(\mathbf{y},\mathbf{y}(n);\mathbf{a})$ falls rapidly for large $|\mathbf{y}-\mathbf{y}(n)|$, as we shall always choose, only members of the data set $\mathbf{y}(n)$ near $\mathbf{y}$, i.e., the neighbors, play much of a role in its future. Our $\mathbf{F}(\mathbf{y},\mathbf{a})$ in that sense is an analytic formulation of the FS idea. More or less weight can be given to the near neighbors by different choices for the function $g(\mathbf{y},\mathbf{y}(n);\mathbf{a})$. The Gaussian we work with could be replaced by a Lorentzian or other choices which weight neighbors more.

C. Invariants

With a map and a cost function, Eqs. (1), (2), and (4), we are ready for the constraints. Section III is devoted to a discussion of Lyapunov exponents. In it we first turn to the extraction of the Lyapunov exponents $\lambda_1,\lambda_2,\ldots,\lambda_d$ from the data $\mathbf{y}(1),\mathbf{y}(2),\ldots,\mathbf{y}(N)$. We do not add anything but our own experience to that of many workers who have explored the calculation of λ_i from data. We attempt to convey to the reader an overview of the available methods for determining Lyapunov spectra and a sense of their reliability. Therefore that portion of Sec. III may be skipped by persons with experience. We include it here since determining the λ_i's is an essential step in our plan for determining $\mathbf{F}(\mathbf{y},\mathbf{a})$ and we have chosen to comment on how we have done it rather than refer the reader to the literature. Of course, we do that too. That established, we discuss how to determine these numbers in terms of the $\mathbf{F}(\mathbf{y},\mathbf{a})$. Equating the numerical values for λ_i from the data to their expression in terms of parameters $\mathbf{a}$ in $\mathbf{F}(\mathbf{y},\mathbf{a})$ will constitute our first set of constraints on the minimization of $C(\mathbf{X},\mathbf{a})$.

Section IV contains our discussion of the invariant distribution of points on the attractor. In principle, this quantity, which we called $\rho(\mathbf{y})$, contains an infinite amount of information on the dynamics. A finite data set $\mathbf{y}(n)$ restricts the resolution we have of this information. We have chosen to express this finite amount of information in terms of the projection of $\rho(\mathbf{y})$ on a set of dual basis functions which are a complete set in $\mathbf{R}^d$. Keeping a finite number of these functions is equivalent to a finite resolution view of the complex structure of $\rho(\mathbf{y})$.[19]

One of our contributions in this work is a scheme for choosing the dual basis functions "tuned" to the structure of $\rho(\mathbf{y})$. This allows us to represent our finite resolution of $\rho(\mathbf{y})$ by a small number of terms in an expansion in the *optimal basis functions*.[25–27] By projecting the $\rho(\mathbf{y})$ determined from the data onto these basis functions, we can determine the coefficients of the expansion of $\rho(\mathbf{y})$ in this basis. Similarly, we can project the $\rho(\mathbf{y})$ determined from the map $\mathbf{F}(\mathbf{y},\mathbf{a})$ onto these basis functions and determine the expansion coefficients of the map. Equating the coefficients one determines from the data to the ones determined from the map constitutes our final constraints on the minimization of $C(\mathbf{X},\mathbf{a})$. Furthermore, we show how the components of $\rho(\mathbf{y})$, in this basis, are the elements of the eigenvalue unity eigenvector of a finite-dimensional matrix constructed from $\mathbf{F}(\mathbf{y},\mathbf{a})$ and the dual basis functions.

In Sec. V we describe our implementation of the constrained minimization program[28] for two model systems: the Hénon map of the plane to itself and (2) the Lorenz system. In each case we numerically generate a data set of $x(n)$'s. We then discuss in some detail our experience in establishing the dimension of the space in which the dynamics is embedded. We also discuss the calculation of Lyapunov exponents, and aspects of the invariant distribution on the attractor from the $\mathbf{y}(n)$'s. Finally, we carry out the constrained minimization of the cost function and indicate how well our parametrized mappings are able to perform in predicting orbits other than those in the given data set.

In this paper we are attempting to describe a method of analyzing *experimental* data. For such a situation we do not know *a priori* the correct embedding dimension, the correct Lyapunov exponents, or the underlying dynamical system that can be used to generate the correct invariant distribution. Yet we have used data sets generated by a dynamical system that we know. We have used known systems for two reasons. The first is that it provides a simple way to obtain large, noise free, data sets. Second, it provides a way of measuring how well existing techniques are able to determine the embedding dimension and the Lyapunov exponents. In order to simulate experimental systems we treat the data set as having come to us from an unknown source. Thus we do not use any of the known properties of either the Hénon or the Lorenz system.

An issue of some importance we do not address in this paper is that of extrinsic *noise* which could contaminate

our signal $x(n)$. This is not a dismissal of this important issue but an attempt to separate out the matters of efficiency and utility of our plan for prediction on strange attractors from issues concerning the practical degradation of our procedures by external noise. An equally important issue is the quantity of data available. The determination of Lyapunov exponents is very difficult for short data sets. As we have stated above the resolution of $\rho(\mathbf{y})$ is determined by the number of data points available. As the dimension of the phase space increases, the amount of data necessary for accurate prediction increases dramatically. We will return to the implications of noisy and/or short data sets for our prediction procedure in later work. For now we assume that we are given essentially noise-free, arbitrarily long time series $x(n)$.

It is our expectation that our experiences with the two systems listed above will give us the ability, in many instances, to construct models in the form of our parametrized mapping $\mathbf{F}(\mathbf{y},\mathbf{a})$ which allow prediction and control of the underlying nonlinear dynamical system producing an observed signal $x(n)$. The details of the $\mathbf{F}(\mathbf{y},\mathbf{a})$ for a specific application should reflect the known features of the physical or other phenomena giving the signal. It seems too bold, if at all possible, to suggest any general rules for choosing forms for $\mathbf{F}(\mathbf{y},\mathbf{a})$. This is sure to be a rich area for experimentation and our own choice will be motivated by considerations we shall defend in a later section and slightly alluded to above.

The matter of noise will be addressed in a future paper. Our methods for dealing with noise follow those outlined by Fuller[29] and seem similar to the ideas of Sidorwich.[30]

II. CHOICE OF AN EMBEDDING SPACE

In this section we illustrate how one can determine the phase-space embedding dimension d from the scalar time series $x(n)$, $n=1,\ldots,N_D$. We assume that the data set is long enough that we need not be concerned with statistical issues about the numerical accuracy of the quantities we consider below. We also assume extrinsic noise is absent from the $x(n)$'s when we receive them. Matters of short and/or noisy data sets, while critical in all applications, are addressed only peripherally in this paper.

Following the work of Packard *et al.*[5] and Mañé and Takens[6,7] and the developmental work of numerous others we seek a set of lagged variables $x(n), x(n+\tau_1), x(n+\tau_2),\ldots,x(n+\tau_{d-1})$ which act as the coordinates in a d-dimensional space in which the dynamics producing the $x(n)$'s is fully captured or embedded.

The choice of lags τ_a is not a well agreed upon matter.[31] The issue is the accuracy and efficiency with which the d-dimensional vectors that result from a particular choice of τ_a's represents the phase space in which the attractor resides. If the underlying system were a differential equation and a scalar variable $x(t)$ were measured at discrete times $x(n)=x(t_0+n\Delta t)$, then we are by the choice of lagged variables trying to find a discrete replacement for the usual phase-space coordinates $x(t), dx/dt,\ldots,d^{d-1}x/dt^{d-1}$. Mañé and Taken's results indicate that, in principle, any choice of lags τ_a will do. We adopt the practice of choosing a single lag τ and making all other lags multiples of τ. The question of what is the best way to choose τ is still open. In a heuristic sense, if τ is too small, then the coordinate at $x(n+\tau)$ and $x(n+2\tau)$ represent almost the same information. Similarly, if τ is too large, then $x(n+\tau)$ and $x(n+2\tau)$ represent distinct uncorrelated descriptions of the embedding space.

For reasons of consistency and ease in calculating Lyapunov exponents (cf. Sec. III) we adopt the following practice. We take the original scalar measurements and calculate its autocorrelation function

$$\frac{1}{T}\int_0^T x(t+\tau)x(t)dt\ .$$

We then choose τ to be approximately $\frac{1}{10}$ to $\frac{1}{20}$ the time associated with the first local minimum of the autocorrelation function. We find that this system, although somewhat arbitrary, works well in practice and provides a simple and systematic way of determining τ. We set τ to unity and thereby establish a time scale for the problem. The data $x(n)$, $n=1,\ldots,N_D$ thus become measurements of the scalar variable separated by a constant time step τ.

We then form d vectors

$$\mathbf{y}(n)=(x(n),x(n+1),\ldots,x(n+d-1)) \tag{5}$$

for $n=1,2,\ldots,\ N=N_D-d$ in a space $\mathbf{R}^d$ capturing the geometric structure of the attractor on which the orbits $x(n)$ lie. To establish d we need some characteristic of the attractor that becomes unchanging as d becomes large enough, thus indicating that the attractor can be embedded in $\mathbf{R}^d$. The usual Hausdorff or other *dimensions* of the attractor are such characteristic quantities. Numerical calculations of the Hausdorff dimension $d_A(N,d)$ of an attractor may depend on the finite length of the data set N and/or the embedding dimension d. For N large enough d_A will become independent of d when the attractor is properly embedded in $\mathbf{R}^d$. Operationally one increases d until d_A remains constant and identifies the minimum d where d_A "saturates" as the embedding dimension.

In fact, we, along with numerous others, do not recommend the computation of d_A, however geometrically appealing it may be, because it is too demanding of computer time. We suggest, and we use, the properties of the correlation function $D(r)$, proposed by Takens[32] and by Grassberger and Procaccia,[33] which is much easier to compute.[34] In terms of the data vectors $\mathbf{y}(n)$ this is defined to be

$$D(r,N,d)=\frac{2}{N(N-1)}\sum_{i=1}^{N}\sum_{j=1}^{N}\Theta(r-|\mathbf{y}(j)-\mathbf{y}(i)|)\ ,\qquad i\neq j \tag{6}$$

where $\Theta(x)$ is the Heaviside function $\Theta(x>0)=1$ and $\Theta(x<0)=0$. The vertical bars represent some measure of distance in $\mathbf{R}^d$—we use the Euclidean norm, but that is only a convenient choice. This correlation function counts the points of the attractor within a distance r of each other. Thus it possesses much of the same geometri-

cal content as the Hausdorff or other invariant dimension attributes.

For N large enough the behavior of $D(r,N,d)$ for small r becomes independent of N. As one would expect from scaling arguments about fractals, as well as observationally, $D(r,N,d)$ is seen to take the form

$$D(r,N,d)=\Phi(r,d)r^{\nu(d)}$$

for small r and large N.[35]

We will identify as the embedding dimension that value of d where the structure in $D(r,N,d)$ becomes independent of d. In this regime it is sufficient that $D(r,N,d)$, becomes independent of d over a range of r near $r\to 0$, and large N [$r=0$ in a finite data set always gives strictly zero for $D(r,N,d)$ and is an uninteresting limit].

To illustrate the use of the correlation function as an embedding dimension discriminant we have calculated $D(r,N,d)$ for very long time series taken from the two examples we will be working with in this paper: (i) the Hénon map of the plane to itself,[36]

$$\begin{aligned} x_1(n+1)&=1.0-ax_1(n)^2+x_2(n)\,,\\ x_2(n+1)&=bx_1(n)\,, \end{aligned} \tag{7}$$

with conventional parameter values $a=1.4$ and $b=0.3$, and (ii) the Lorenz system of three autonomous differential equations[37]

$$\begin{aligned} \frac{dx_1(t)}{dt}&=\sigma(x_2(t)-x_1(t))\,,\\ \frac{dx_2(t)}{dt}&=-x_1(t)x_3(t)+rx_1(t)-x_2(t)\,,\\ \frac{dx_3(t)}{dt}&=x_1(t)x_2(t)-bx_3(t)\,, \end{aligned} \tag{8}$$

with parameter values $\sigma=16$, $b=4$, and $r=45.92$.

For the Hénon map we took an initial condition lying in its basin of attraction and iterated the map 4550 times. The first 50 iterates were discarded as representing transient behavior, while the last 4500 points of $x_1(n)$ and $x_2(n)$ were then used to make d vectors

$$\mathbf{y}_i(n)=(x_i(n),x_i(n+1),\ldots,x_i(n+d-1))$$

for $i=1$ or 2. For $d=1,2,\ldots,5$ $D(r,N,d)$ was computed using an efficient code developed by Theiler.[34] For $\mathbf{y}_1(n)$ data these $D(r,N,d)$ are plotted in Fig. 1. A similar plot was generated for $\mathbf{y}_2(n)$, but is not shown. Because of the simplicity of the connection $x_2(n+1)=bx_1(n)$ in the Hénon map, these two views of $D(r,N,d)$ are really redundant. However, in the spirit of treating each data series as having originally come to us from a source whose underlying dynamics is unknown we performed both calculations.

While a cautious and careful observer might say the embedding dimension for the $\mathbf{y}_1(n)$ data would be $d=3$, we feel safe in concluding from these figures that $d=2$. Computations with N greater than 4500 support this conclusion.

Further, if we take the $x_1(n)$ data and plot the two-

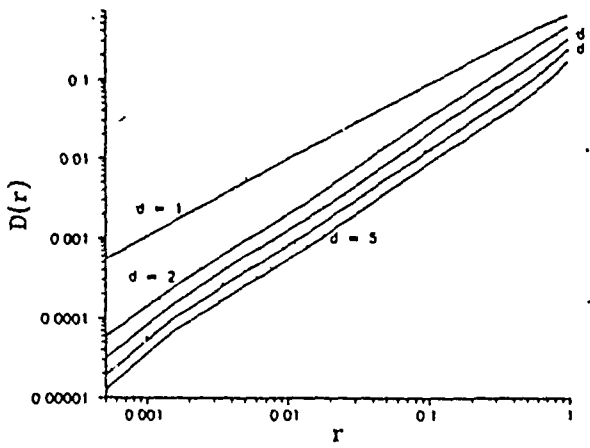

FIG. 1. $D(R)$ vs r for the Hénon map. $\mathbf{y}_1(n)=(x_1(n), x_1(n+1))$ for 4500 points.

vectors $\mathbf{y}_1(n)=(x_1(n),x_1(n+1))$, we reconstruct the figure seen in Fig. 2. This is, as should not be surprising in this simple example, a rotated form of the Hénon attractor. The usual display of the Hénon attractor is obtained by plotting $(x_1(n),x_2(n))$ for our data. Since $x_1(n)$ is $(1/b)x_2(n+1)$, the coincidence of these plots is certainly not remarkable. Our goal in presenting this kind of detail is as a guide to what one might expect in more complicated examples rather than as revelations about the Hénon map.

Next we turn to the Lorenz equations. Once again we chose initial conditions in the basin of attraction and solved Eqs. (8) with a straightforward fourth-order Runge-Kutta ordinary differential equation (ODE) solver with a fixed time step. After discarding the first 50 time steps as transients, we recorded x_1, x_2, and x_3 for $N=4500$ corresponding to many natural cycles of the orbit around the attractor. From each of the three data sets we formed the d vectors as in Eq. (5) and with them evaluated the correlation function $D(r,N,d)$ for $d=1,2,\ldots,5$. The $D(r,N,d)$'s for $\mathbf{y}_1(n)$ data are shown

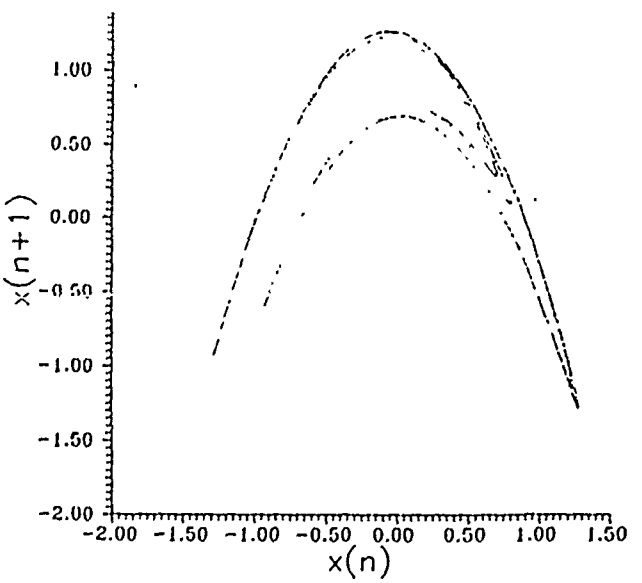

FIG. 2. Hénon attractor $x_1(n)$ plotted against $x_1(n+1)$.

in Fig. 3. An embedding dimension of $d=3$ is fairly clearly a safe choice for these data. A bolder choice would have been $d=2$. Since it is known that the Hausdorff dimension of the Lorenz attractor is just above 2 in this regime of parameter space, this would have been a convincing, although incorrect, choice. The message here is that choosing d too large entails extra subsequent computation, but no loss of information on the attractor. It is probably safer to live with a d one dimension too large as a general matter of care. We thus choose $d=3$. The results of the y_2 and y_3 data are not shown. As with the Hénon example the results of y_2 and y_3 are similar to those of y_1. The fact that the y_1, y_2, and y_3 vectors for the Lorenz data (y_1 and y_2 for Hénon) yields similar results is to be expected since all three measurements evolve on the same attractor.

Next we plot the points $y_1(n)=(x_1(n),x_1(n+1),x_1(n+2))$ in the three-dimensional embedding space. These are shown in Fig. 4 as a projection on a plane with normal vector $\hat{n}=(\cos\theta)\hat{x}_1(n)+(\sin\theta)\hat{x}_1(n+1)+0\hat{x}_1(n+2)$ for $\theta=1.31$. We note the similarity between Fig. 4 and the well-known shape of the Lorenz attractor. Thus the method of *phase-space embedding* reliably reproduces the Lorenz attractor. For the two examples we have used the reconstructed attractor is similar in appearance to the attractor generated by the "true" underlying equations of motion. In general, the reconstructed attractor will not have this visual similarity. However, the reconstructed attractor will contain all of the important invariant information as the true attractor. The difference in visual shapes is the result of a nonlinear change of variables between the true dynamical variables and the reconstructed variables.

We close this section with a summary note reminding the reader that our use of the correlation integral Eq. (6) has been to establish an embedding dimension d in which to view the system attractor described by our time series $x(n)$. We chose $D(r,N,d)$ because it is familiar, easy to compute, and has a clear geometrical meaning. For us it is a diagnostic tool. While the details of the small r behavior $D(r,d)\sim r^{\nu}\Phi(r)$ contains important information about the dynamics, we do not focus on that here. Indeed, we are quite happy to accept other diagnostic tools in its place.

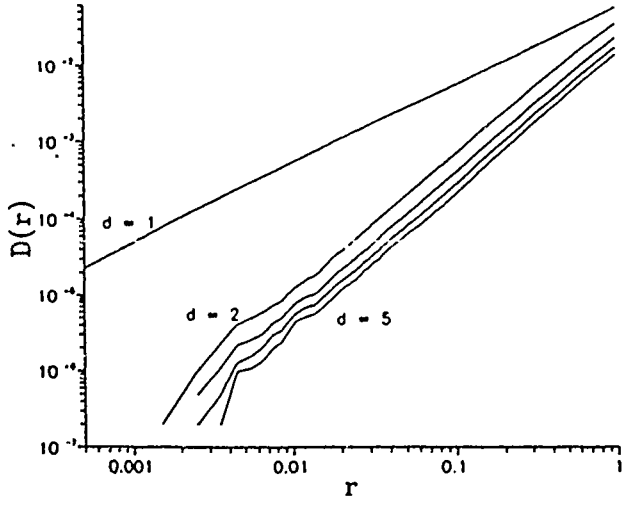

FIG. 3. $D(r)$ vs r for the Lorenz equations, $y_1(n)$ for 4500 points and embedding dimensions $d=1,\ldots5$. For this case $r=45.92$, $b=4.0$, and $\sigma=16$.

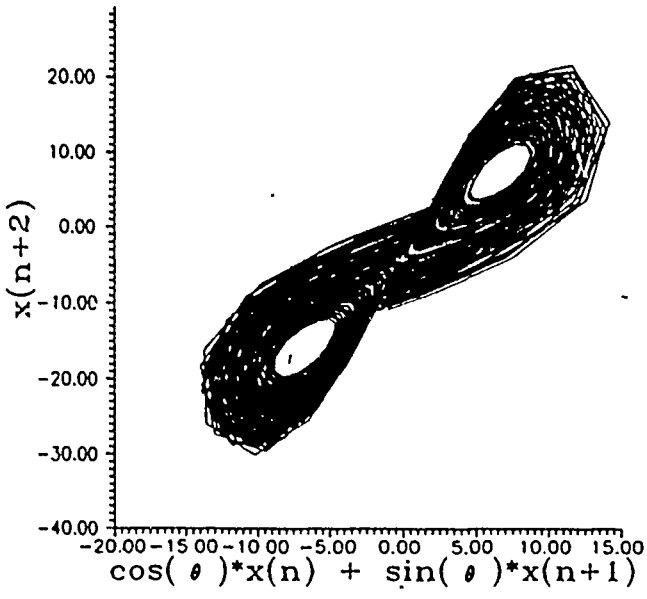

FIG. 4. Lorenz attractor created from $x_1(n)$ data. The parameter values are $r=45.92$, $b=4.0$, and $\sigma=16$, while the projection angle is $\theta=1.31$.

III. LYAPUNOV CHARACTERISTIC EXPONENTS—FROM DATA AND FROM THE MAP

In this section we discuss how one determines the Lyapunov exponents that govern a dynamical system. First we discuss how to extract them from an experimental data set and then from our mapping $\mathbf{F}(\mathbf{y},\mathbf{a})$. By choosing the parameters $\mathbf{a}$ in such a way that $\mathbf{F}(\mathbf{y},\mathbf{a})$ yields the same Lyapunov exponents as the experimental data set, we are forcing a constraint on $\mathbf{F}(\mathbf{y},\mathbf{a})$ that is not explicitly required by minimizing the cost function given by Eq. (4). This local constraint should improve our ability to predict the short-term (and possibly long-term) evolution of points that are not in the data set, but near the attractor. Certainly points outside the basin of attraction of the attractor we have observed in the original data set will not evolve according to our $\mathbf{F}(\mathbf{y},\mathbf{a})$.

Rather than writing our own computer program, and thereby become embroiled in the controversy of what is the best way to determine Lyapunov spectra from an experimental time series, we have chosen to use methods that have already been proposed by two different research groups. By comparing the results of both methods we hope to improve our confidence in the spectra given by each of them separately. The first method we shall report on was developed by Eckmann *et al.*[38] The second method was developed by Wolf *et al.*[39] Finally, we will show how we calculated the Lyapunov spectra from our mapping $\mathbf{F}(\mathbf{y},\mathbf{a})$.

The choice of an appropriate data set for use in either the Eckmann *et al.* or the Wolf *et al.* method is some-

thing that cannot be overstressed. As stated in Sec. II the time lag τ between successive measurements of the dynamical variable must be appropriately chosen, if one wants optimal results.

A. Eckmann-Kamphorst-Ruelle-Cilliberto method

For the Eckmann *et al.* method the FORTRAN source code we used when performing our numerical experiments on the dynamical systems denoted in Sec. II was provided by the authors of Ref. 38. It assumes that the input is a string of positive integer data whose sampling rate is τ. [The temporary conversion of the data set $x(n)$, $n=1,\ldots,N_D$ to positive integers for the sake of the Lyapunov calculation should not be a difficult matter.] The code reads the data set and embeds it in a d-dimensional space in the manner specified in Sec. II. The result is a set of $N=N_D$ data vectors $\mathbf{y}(n)=(x(n),x(n+1),\ldots,x(n+d-1))$ where we have normalized τ to unity [cf. Eq. (5)]. It then chooses an initial $\mathbf{y}(n)$ and finds all neighbors of $\mathbf{y}(n)$ within a radius r. These points, as well as their forward images, are used to construct a linear mapping T from time n to time $n+1$. The Lyapunov exponents are related to the eigenvalues of the successive iterates of the map T. For a detailed discussion of the algorithm we direct the reader to Ref. 38.

The Eckmann *et al.* method assumes that the embedding dimension d is related to the number of Lyapunov exponents via the rule $d=(d_m-1)M+1$, where d_m is the number of Lyapunov exponents and M is a positive integer. By allowing d_m and M to range over various values a wide range of embedding dimensions is used. We remark that the reader will recall that in Sec. II we established a method for determining the minimum embedding dimension d. The data vectors $\mathbf{y}(n)$ are assumed to live on some attractor that occupies some portion of $\mathbf{R}^d$. It is a numerically difficult exercise to calculate Lyapunov exponents from data. Thus it is necessary to examine a wide range of possible embedding dimensions d. It is our experience that the calculated values of the exponents converge onto their correct values as d is increased above the number specified by methods in Sec. II. We report numerical experiments for d_m in the range between 2 and 9 for $M=1,2$. (We remark that $M=1$ recovers $d=d_m$, while $M=2$ is slightly below the Takens and Mañé limit.[6,7]) In all of our tests we iterated the tangent map T 2000 times before evaluating the Lyapunov exponent.

To get a feel for the proper densities of points on the reconstructed attractor, it is useful to use diagnostics such as, say, a histogram of the number of neighbors falling within a range around the smallest nearest-neighbor distance on the attractor. If the density of points on an attractor is quite inhomogeneous, much higher mean point densities are often necessary to insure that most points have at least a few nearby neighbors. Often a useful diagnostic is simply to plot out the reconstructed attractor, and visually obtain an intuitive feel for how homogeneous the point density is. As a general rule of thumb (inspired by Wolf *et al.*), we find empirically that the minimum number of points required for the prediction algorithm to go as something like 20^d, where d is the dimension of the embedding space, although this is probably an overestimate when d is 4 or more.

The first dynamical system for which we present results is the Hénon map of the plane given by Eqs. (7). We used a data set with $N=10\,000$ entries. The results are shown in Fig. 5 and Table I. As one can see, the numerical experiments accurately predict the accepted value of the positive Lyapunov exponent $\lambda_1=0.418$. Although for the $M=1$ case the computer code produced a reasonable accurate prediction of the negative Lyapunov exponent, the code, in principle, will not yield accurate values of the negative or zero Lyapunov exponents. This fact is born out in the $M=2$ case (which is not shown). Furthermore, we know of no method that will produce negative Lyapunov exponents from an experimental data set. Since we are unable to reliably determine the negative Lyapunov exponents from the data, we will not constrain $\mathbf{F}(\mathbf{y},\mathbf{a})$ to reproduce the negative values of the spectra.

It should not be surprising that we are unable to determine the negative Lyapunov spectra using our data sets. We have assumed that the data describe motion *on* an attractor. The negative Lyapunov exponents indicate how points in the phase space that are *off* the attractor get onto the attractor. The portion of the data set that might reveal how points off the attractor get to the attractor is the initial transient. This transient is typically very short (sometimes as few as 10 time steps τ) and is usually discarded or otherwise unavailable.

A related issue to be addressed is that the code produces d_m exponents regardless of the actual number of Lyapunov exponents that govern the dynamics of the

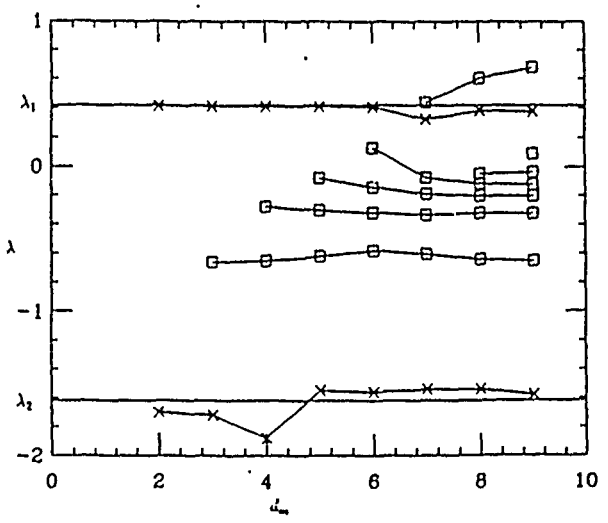

FIG. 5. The results of calculating Lyapunov exponents by the Eckmann *et al.* method for Hénon data. The horizontal axis is d_m, the assumed dimension of the dynamical system that produced the data set. Thus the method will produce d_m Lyapunov exponents. The vertical axis contains the numerical values calculated for the d_m different λ's. The two horizontal lines are the known correct values for $\lambda_1=0.418$ and $\lambda=-1.62$. The method relates d_m to the embedding dimension d via $d=(d_m-1)M+1$. This figure shows results for $M=1$. Spurious exponents are labeled with squares while dynamical exponents are labeled with X's.

TABLE I. Lyapunov exponents for the Hénon attractor $M=1$, and the number of data points is 10000.

d_m	λ				
2	$\lambda_1=0.412$	$\lambda_2=-1.70$			
3	$\lambda_1=0.412$	$\lambda_2=-0.662$	$\lambda_3=-1.72$		
4	$\lambda_1=0.408$	$\lambda_2=-0.281$	$\lambda_3=-0.655$	$\lambda_4=-1.88$	
5	$\lambda_1=0.408$	$\lambda_2=-0.0824$	$\lambda_3=0.305$	$\lambda_4=-0.622$	$\lambda_5=-1.55$
6	$\lambda_1=0.407$	$\lambda_2=0.128$	$\lambda_3=-0.144$	$\lambda_4=-0.321$	$\lambda_5=-0.581$
	$\lambda_6=-1.56$				
7	$\lambda_1=0.437$	$\lambda_2=0.323$	$\lambda_3=-0.0767$	$\lambda_4=-0.190$	$\lambda_5=-0.335$
	$\lambda_6=-0.604$	$\lambda_7=-1.54$			
8	$\lambda_1=0.602$	$\lambda_2=0.382$	$\lambda_3=-0.0509$	$\lambda_4=-0.118$	$\lambda_5=-0.203$
	$\lambda_6=-0.332$	$\lambda_7=-0.642$	$\lambda_8=-1.54$		
9	$\lambda_1=0.677$	$\lambda_2=0.377$	$\lambda_3=0.0896$	$\lambda_4=-0.0390$	$\lambda_5=-0.124$
	$\lambda_6=-0.203$	$\lambda_7=-0.324$	$\lambda_8=-0.652$	$\lambda_9=-1.58$	
Accepted values of λ		$\lambda_1=0.418$		$\lambda_2=-1.62$	

physical system in question. However, it is relatively easy to determine which of the d_m exponents govern the dynamics of the system and which are spurious. We assume that one has successfully determined the minimum embedding dimension of the attractor by the method we presented in Sec. II (or any other reliable method at the reader's disposal). Examination of Fig. 5 indicates that most of the spurious exponents are negative. These negative exponents are necessary to contract the d-dimensional phase space onto the attractor whose dimension is $d_A<d$. The one positive spurious exponent appears at $d_m=7$ for the $M=1$ case. We know from Fig. 1 that the dynamics of the Hénon attractor can be embedded in two dimensions. Hence we conclude that an exponent that exist only for $d_m\geq 7$ must be spurious. The origin of this spurious positive exponent is discussed by Eckmann *et al.*[38] It is believed it will stabilize at a value of $2\lambda_1$.

We have averaged the calculated values of λ_1 for the $M=1$ case over the range $d_m=2$–6. We discarded the values of λ_1 for $d_m\geq 7$ since they have obviously been altered by the spurious Lyapunov exponent generated at $d_m=7$. We find that the average value is $\bar{\lambda}_1=0.409$, which differs from the accepted value of 0.418 by only 2%. For the $M=2$ case we found similar results. After averaging we find that $\bar{\lambda}_1=0.420$. In conclusion, we state that by comparing the $M=1$ and 2 cases we feel that the code successfully determined the positive Lyapunov exponent associated with the Hénon attractor.

We now turn our attention to the second dynamical system we wish to analyze, the Lorenz system of ODE's given by Eqs. (8). The data set used for our numerical experiments consisted of $N=20000$ entries and was generated by integrating Eqs. (8) forward in time using a simple fourth-order Runge-Kutta routine with a fixed time step. We chose to record the $x_1(t)$ variable, although either the $x_2(t)$ or $x_3(t)$ variable would do as well. Figure 6 is a graph of the autocorrelation function. The first minimum is at $n\sim 12$ where n is the number of Runge-Kutta time steps of length 0.03. The time associated with this first minimum is approximately $t_c\sim 0.36$. We use a sampling rate $\tau=0.03$, which is approximately $\frac{1}{12}$ of the autocorrelation time. Thus we use every Runge-Kutta data point as our experimental data set. We allowed d_m to range between 2 and 9 for $M=1$ and 2. The results of our numerical experiments are shown in Fig. 7 and Table II.

For all cases $M=1$ and 2, we are able to accurately determine both the positive and the zero Lyapunov exponent. The accepted value of λ_1 is 1.50. The average of the calculated values of λ_1 for $d_m\geq 5$ in the $M=1$ case is $\bar{\lambda}_1=1.45$, which is an error of only 3%. As with the Hénon example, we found better results for the $M=2$ case.

The question of a zero Lyapunov exponent requires special consideration. Any dynamical system that is represented by an ODE will contain a zero Lyapunov exponent. As can be seen from Fig. 7 and Table II, one of the Lyapunov exponents calculated from the experimental data set is very small (as much as two orders of magnitude) compared to λ_1. We also notice that this exponent is very stable and very persistent. It exists for $M=1$ and 2 over the entire range of d_m. Given this behavior and

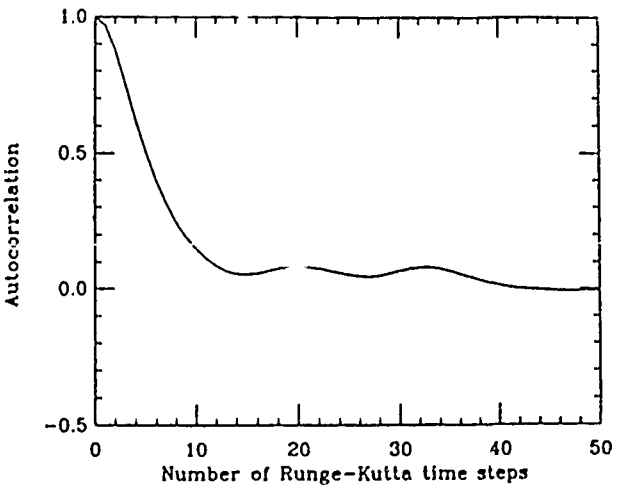

FIG. 6. Autocorrelation function for $x_1(t)$, the Lorenz system from 20000 data points.

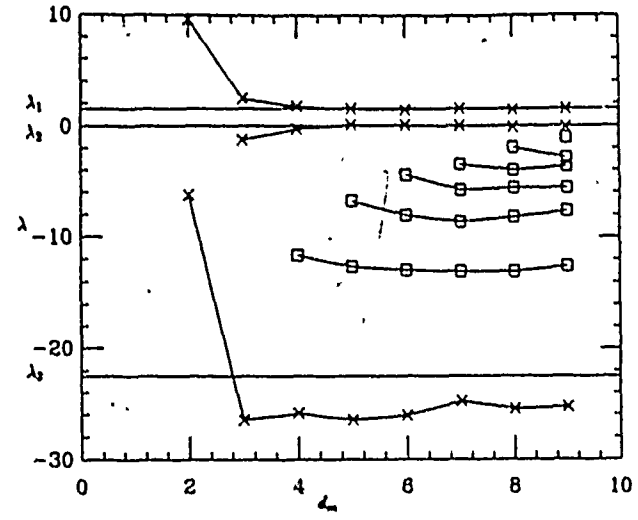

FIG. 7. The results of applying the Eckmann *et al.* method of calculating Lyapunov exponents to a Lorenz data set. The results shown are for $M=1$. For this dynamical system there are three dynamical exponents at $\lambda_1=1.50$, $\lambda_2=0.0$, and $\lambda_3=-22.5$.

the prevalence of ODE's as dynamical systems, we feel confident in predicting a zero Lyapunov exponent.

Of course, we have the luxury here of *knowing* that our data set came from an ODE. This type of knowledge concerning the origin of a data set is typically unavailable. Thus we must use our best judgment and live with the fact that we cannot know for *certain* whether a Lyapunov exponent generated by the Eckmann *et al.* method should be interpreted as zero or just very small. Our recommendation is that one compare the suspected zero exponent to the smallest nonspurious positive Lyapunov exponent generated by the code. If the suspected exponent is as persistent, as stable, and more than a factor of 25 smaller than the smallest positive exponent, we recommend that the suspected exponent be assigned the value zero.

B. Wolf-Swift-Swinney-Vastano method

A second technique that we have investigated to determine Lyapunov exponents from time series is due to Wolf, Swift, Swinney, and Vastano (WSSV).[39] This paper presents two algorithms, one for determining the full Lyapunov spectrum from a known set of dynamical equations, and one for determining only the largest positive exponent if one has available only a time series from the dynamical system. Since the paper includes the source codes for the two algorithms, we copied and used them directly. The WSSV code for time series analysis can only determine the largest positive exponent. We have up to now chosen to use only one Lyapunov exponent as a constraint to the nonlinear fitting method, and so this program proves sufficient for our needs. Given the current difficulty of accurately determining other exponents from a time series of data, we restrict the constraints to one Lyapunov exponent. In addition to these considerations, the WSSV code is exceptionally easy to use, and requires relatively minimal amounts of data.

The WSSV code for time series works in a manner somewhat similar to other techniques which attempt to approximate in some way the local tangent space about a fiducial orbit. In this case, the code initially constructs the time-delay reconstructed coordinates for the system in the usual manner, taking the parameters for the reconstruction as input to the program. The calculation of the Lyapunov exponent then begins by finding the nearest neighbor in the reconstructed phase space to the first point of the orbit, where "nearest" is measured using the usual Euclidean metric. Once this point is found, the magnitude of the difference vector between the two points is recorded. The algorithm then proceeds by evolving the fiducial point along its trajectory, and the neighboring point along its trajectory, a given number of steps of the time series. The magnitude of the final separation between the two points is then determined, and the contribution to the Lyapunov exponent is then simply given as the logarithm of the final separation divided by the initial separation, divided by the time interval of evo-

TABLE II. Lyapunov exponents for the Lorenz attractor $M=1$, and the number of data points is 20 000.

d_m	λ				
2	$\lambda_1=9.54$	$\lambda_2=-6.30$			
3	$\lambda_1=2.42$	$\lambda_2=-1.27$	$\lambda_3=-26.5$		
4	$\lambda_1=1.68$	$\lambda_2=-0.308$	$\lambda_3=-11.7$	$\lambda_4=-25.9$	
5	$\lambda_1=1.47$	$\lambda_2=0.0619$	$\lambda_3=-6.84$	$\lambda_4=-12.7$	$\lambda_5=-26.5$
6	$\lambda_1=1.40$	$\lambda_2=0.0471$	$\lambda_3=-4.50$	$\lambda_4=-8.12$	$\lambda_5=-13.0$
	$\lambda_6=-26.1$				
7	$\lambda_1=1.50$	$\lambda_2=0.0141$	$\lambda_3=-3.49$	$\lambda_4=-5.81$	$\lambda_5=-8.69$
	$\lambda_6=-13.1$	$\lambda_7=-24.8$			
8	$\lambda_1=1.40$	$\lambda_2=-0.105$	$\lambda_3=-1.96$	$\lambda_4=-4.02$	$\lambda_5=-5.62$
	$\lambda_6=-8.24$	$\lambda_7=-13.1$	$\lambda_8=-25.5$		
9	$\lambda_1=1.48$	$\lambda_2=-0.109$	$\lambda_3=-1.06$	$\lambda_4=-2.88$	$\lambda_5=-3.68$
	$\lambda_6=-5.63$	$\lambda_7=-7.70$	$\lambda_8=-12.6$	$\lambda_9=-25.3$	
	Accepted values of λ		$\lambda_1=1.50$	$\lambda_2=0.00$	$\lambda_3=-22.5$

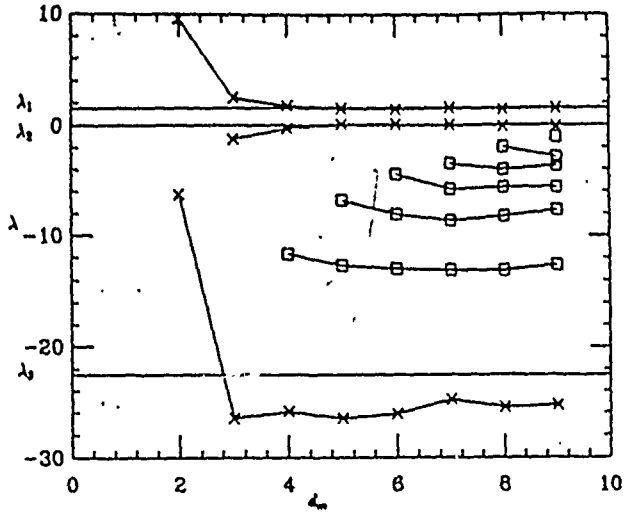

FIG. 7. The results of applying the Eckmann *et al.* method of calculating Lyapunov exponents to a Lorenz data set. The results shown are for $M=1$. For this dynamical system there are three dynamical exponents at $\lambda_1=1.50$, $\lambda_2=0.0$, and $\lambda_3=-22.5$.

the prevalence of ODE's as dynamical systems, we feel confident in predicting a zero Lyapunov exponent.

Of course, we have the luxury here of *knowing* that our data set came from an ODE. This type of knowledge concerning the origin of a data set is typically unavailable. Thus we must use our best judgment and live with the fact that we cannot know for *certain* whether a Lyapunov exponent generated by the Eckmann *et al.* method should be interpreted as zero or just very small. Our recommendation is that one compare the suspected zero exponent to the smallest nonspurious positive Lyapunov exponent generated by the code. If the suspected exponent is as persistent, as stable, and more than a factor of 25 smaller than the smallest positive exponent, we recommend that the suspected exponent be assigned the value zero.

B. Wolf-Swift-Swinney-Vastano method

A second technique that we have investigated to determine Lyapunov exponents from time series is due to Wolf, Swift, Swinney, and Vastano (WSSV).[39] This paper presents two algorithms, one for determining the full Lyapunov spectrum from a known set of dynamical equations, and one for determining only the largest positive exponent if one has available only a time series from the dynamical system. Since the paper includes the source codes for the two algorithms, we copied and used them directly. The WSSV code for time series analysis can only determine the largest positive exponent. We have up to now chosen to use only one Lyapunov exponent as a constraint to the nonlinear fitting method, and so this program proves sufficient for our needs. Given the current difficulty of accurately determining other exponents from a time series of data, we restrict the constraints to one Lyapunov exponent. In addition to these considerations, the WSSV code is exceptionally easy to use, and requires relatively minimal amounts of data.

The WSSV code for time series works in a manner somewhat similar to other techniques which attempt to approximate in some way the local tangent space about a fiducial orbit. In this case, the code initially constructs the time-delay reconstructed coordinates for the system in the usual manner, taking the parameters for the reconstruction as input to the program. The calculation of the Lyapunov exponent then begins by finding the nearest neighbor in the reconstructed phase space to the first point of the orbit, where "nearest" is measured using the usual Euclidean metric. Once this point is found, the magnitude of the difference vector between the two points is recorded. The algorithm then proceeds by evolving the fiducial point along its trajectory, and the neighboring point along its trajectory, a given number of steps of the time series. The magnitude of the final separation between the two points is then determined, and the contribution to the Lyapunov exponent is then simply given as the logarithm of the final separation divided by the initial separation, divided by the time interval of evo-

TABLE II. Lyapunov exponents for the Lorenz attractor $M=1$, and the number of data points is 20 000.

d_m	λ				
2	$\lambda_1=9.54$	$\lambda_2=-6.30$			
3	$\lambda_1=2.42$	$\lambda_2=-1.27$	$\lambda_3=-26.5$		
4	$\lambda_1=1.68$	$\lambda_2=-0.308$	$\lambda_3=-11.7$	$\lambda_4=-25.9$	
5	$\lambda_1=1.47$	$\lambda_2=0.0619$	$\lambda_3=-6.84$	$\lambda_4=-12.7$	$\lambda_5=-26.5$
6	$\lambda_1=1.40$	$\lambda_2=0.0471$	$\lambda_3=-4.50$	$\lambda_4=-8.12$	$\lambda_5=-13.0$
	$\lambda_6=-26.1$				
7	$\lambda_1=1.50$	$\lambda_2=0.0141$	$\lambda_3=-3.49$	$\lambda_4=-5.81$	$\lambda_5=-8.69$
	$\lambda_6=-13.1$	$\lambda_7=-24.8$			
8	$\lambda_1=1.40$	$\lambda_2=-0.105$	$\lambda_3=-1.96$	$\lambda_4=-4.02$	$\lambda_5=-5.62$
	$\lambda_6=-8,24$	$\lambda_7=-13.1$	$\lambda_8=-25.5$		
9	$\lambda_1=1.48$	$\lambda_2=-0.109$	$\lambda_3=-1.06$	$\lambda_4=-2.88$	$\lambda_5=-3.68$
	$\lambda_6=-5.63$	$\lambda_7=-7.70$	$\lambda_8=-12.6$	$\lambda_9=-25.3$	
Accepted values of λ			$\lambda_1=1.50$	$\lambda_2=0.00$	$\lambda_3=-22.5$

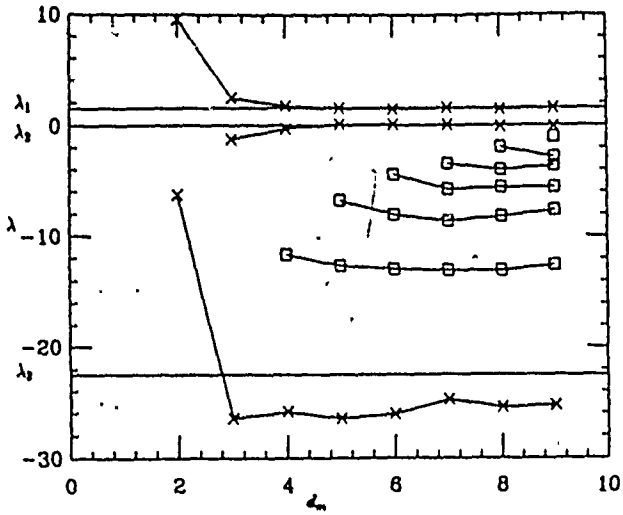

FIG. 7. The results of applying the Eckmann *et al.* method of calculating Lyapunov exponents to a Lorenz data set. The results shown are for $M=1$. For this dynamical system there are three dynamical exponents at $\lambda_1=1.50$, $\lambda_2=0.0$, and

...e r... dynamical systems, we feel ...on... ...a ze... Lyapunov exponent.

Of course, we have the luxury here of *knowing* that our data set came from an ODE. This type of knowledge concerning the origin of a data set is typically unavailable. Thus we must use our best judgment and live with the fact that we cannot know for *certain* whether a Lyapunov exponent generated by the Eckmann *et al.* method should be interpreted as zero or just very small. Our recommendation is that one compare the suspected zero exponent to the smallest nonspurious positive Lyapunov exponent generated by the code. If the suspected exponent is as persistent, as stable, and more than a factor of 25 smaller than the smallest positive exponent, we recommend that the suspected exponent be assigned the value zero.

B. Wolf-Swift-Swinney-Vastano method

A second technique that we have investigated to determine Lyapunov exponents from time series is due to Wolf, Swift, Swinney, and Vastano (WSSV).[39] This paper presents two algorithms, one for determining the full Lyapunov spectrum from a known set of dynamical equations, and one for determining only the largest positive exponent if one has available only a time series from the dynamical system. Since the paper includes the source codes for the two algorithms, we copied and used them directly. The WSSV code for time series analysis can only determine the largest positive exponent. We have up to now chosen to use only one Lyapunov exponent as a constraint to the nonlinear fitting method, and so this program proves sufficient for our needs. Given the current difficulty of accurately determining other exponents from a time series of data, we restrict the constraints to one Lyapunov exponent. In addition to these considerations, the WSSV code is exceptionally easy to use, and requires relatively minimal amounts of data.

The WSSV code for time series works in a manner somewhat similar to other techniques which attempt to approximate in some way the local tangent space about a fiducial orbit. In this case, the code initially constructs the time-delay reconstructed coordinates for the system in the usual manner, taking the parameters for the reconstruction as input to the program. The calculation of the Lyapunov exponent then begins by finding the nearest neighbor in the reconstructed phase space to the first point of the orbit, where "nearest" is measured using the usual Euclidean metric. Once this point is found, the magnitude of the difference vector between the two points is recorded. The algorithm then proceeds by evolving the fiducial point along its trajectory, and the neighboring point along its trajectory, a given number of steps of the time series. The magnitude of the final separation between the two points is then determined, and the contribution to the Lyapunov exponent is then simply given as the logarithm of the final separation divided by the initial separation, divided by the time interval of evo-

TABLE II. Lyapunov exponents for the Lorenz attractor $M=1$, and the number of data points is 20 000.

d_m	λ				
2	$\lambda_1=9.54$	$\lambda_2=-6.30$			
3	$\lambda_1=2.42$	$\lambda_2=-1.27$	$\lambda_3=-26.5$		
4	$\lambda_1=1.68$	$\lambda_2=-0.308$	$\lambda_3=-11.7$	$\lambda_4=-25.9$	
5	$\lambda_1=1.47$	$\lambda_2=0.0619$	$\lambda_3=-6.84$	$\lambda_4=-12.7$	$\lambda_5=-26.5$
6	$\lambda_1=1.40$	$\lambda_2=0.0471$	$\lambda_3=-4.50$	$\lambda_4=-8.12$	$\lambda_5=-13.0$
	$\lambda_6=-26.1$				
7	$\lambda_1=1.50$	$\lambda_2=0.0141$	$\lambda_3=-3.49$	$\lambda_4=-5.81$	$\lambda_5=-8.69$
	$\lambda_6=-13.1$	$\lambda_7=-24.8$			
8	$\lambda_1=1.40$	$\lambda_2=-0.105$	$\lambda_3=-1.96$	$\lambda_4=-4.02$	$\lambda_5=-5.62$
	$\lambda_6=-8.24$	$\lambda_7=-13.1$	$\lambda_8=-25.5$		
9	$\lambda_1=1.48$	$\lambda_2=-0.109$	$\lambda_3=-1.06$	$\lambda_4=-2.88$	$\lambda_5=-3.68$
	$\lambda_6=-5.63$	$\lambda_7=-7.70$	$\lambda_8=-12.6$	$\lambda_9=-25.3$	
	Accepted values of λ		$\lambda_1=1.50$	$\lambda_2=0.00$	$\lambda_3=-22.5$

lution. These contributions are then averaged over the length of the time series.

This simple scheme works to provide the largest Lyapunov exponent because, given arbitrary initial conditions for the two neighbors and an appropriately long evolution time, the exponential growth due to the largest positive exponent dominates the overall behavior of the difference vectors, so that to good accuracy the net change in the magnitude of the two vectors reflects almost solely this rate of growth. Note that a technical problem exists here, in that the Lyapunov exponent is defined only in terms of the linearized equations of motion about the fiducial trajectory, and the exponential divergence of neighboring trajectories can rapidly drive them out of the linear regime. In the WSSV code, this problem is addressed in a straightforward manner; i.e., when the distance between the neighbors becomes too large, the algorithm abandons this point and searches for a new neighbor. A suitable new neighbor is chosen on the basis of two criteria: first, the point must again be as close to the fiducial trajectory point as possible, and second, the orientation of the abandoned difference vector must be preserved as nearly as possible. The process of choosing a new neighbor using these criteria is thus approximately equivalent to rescaling the difference vector to a much smaller size. In practical terms there is a trade-off between choosing points which are very close to the fiducial point and points whose difference vectors lie nearly along the ray defined by the abandoned difference vector. This problem is handled internally in the code by a multistep search algorithm. Once a suitable new neighbor is determined, the new difference vector is then evolved until it too becomes too large, and then the process is repeated.

Because this numerical procedure is relatively straightforward, there are actually few variables necessary as input to the algorithm, and hence the program is much easier to use than other Lyapunov exponent algorithms. There are seven basic variables which must be set to perform an analysis of a data set, most of which are determined when one calculates the embedding dimension as in Sec. II. The first four variables, which are related to the time-delay reconstruction, are the number of points in the data set (N), the embedding dimension d, reconstruction time delay τ_d, and the sampling rate for the data T_s. The first of these variables N is usually fixed when an experimental time series is being analyzed, although some criterion for the minimum number of points necessary for a good estimate of the Lyapunov exponent can be given. Wolf, Swift, Swinney, and Vastano give a general rule for the minimum number of data points as at least 10^d, although this value can depend on the topology of the attractor and the relative magnitudes of the Lyapunov exponents. Our experience has shown that at least twice this number of points is usually necessary for two significant figures of accuracy, and greater accuracy can require much longer time series. It should be noted that in terms of the algorithm, longer time series are required not just to improve the convergence by providing more contributions to the Lyapunov value; longer time series also provide a higher density of points on the attractor and hence there are more nearby neighbors to choose from when replacements are necessary, making this process more accurate.

The embedding dimension parameter d is the dimension of the time-delay reconstructed vectors $\mathbf{y}(n)$, and is determined as in Sec. II. As discussed there, the dimension of the embedding space must be sufficiently large to ensure that none of the dynamical information about the attractor is lost; however, needlessly large values of the embedding dimension results in greatly increased computation time for the Lyapunov calculation and also increased sensitivity to noise. For the example systems that we have investigated using this method, we have chosen the embedding dimension to be the next highest integer dimension to the (known) fractal dimension, although for experimental data, where one is not sure of the fractal dimension, one may often feel safer to choose a larger value.

The second variable that is necessary for the time-delay reconstruction in the program is the actual time delay value τ_d. This variable, as discussed in Sec. II, gives the time separation of the components of the d vectors in terms of the number of iterates of the time series, and can be thought of as being chosen to make the d components as "orthogonal" as possible. For dynamical time series derived from a mapping, as for the Hénon system, this value can be chosen to be 1, since each iterate generally represents one entire "orbit" on the attractor of the flow that the mapping is derived from. For continuous phase-space flows, as for the Lorenz system, one can often use the rule of thumb given by $d\tau_d = 1$, where d is the embedding dimension and τ_d is here given as the fraction of the orbital period on the attractor, which must then be expressed in time series steps. Another more sophisticated method is to take τ_d as roughly the first zero of the autocorrelation function for the time series. The choice of method for determining the time delay is not crucial, however, since the reconstructed dynamics is generally not strongly dependent on the exact value as long as it is within a reasonable range of the correct value.

The fourth variable T_s is the time between successive measurements in the time series, or rather the inverse of the sampling rate. This value is not actually a variable, but rather an additional piece of information that must be supplied with any time series, and is used in the algorithm to rescale the Lyapunov exponents by setting the time scale for the rate of divergence of the trajectories. Although one may have no control over the sampling rate for an arbitrary set, for systems where one does have control this parameter is an important issue, and can greatly affect the quality of data. Many of the aspects of problems that can arise are from improper sampling rates are discussed by Mayer-Kress.[40]

Two of the input variables to the algorithm are concerned with setting length scales for the reconstructed dynamics. The parameter S_{max} controls the maximum distance that the algorithm will look for neighbors when it attempts replacement. Since we take a rough value for the limit of the validity of the linear approximation to be about 1% of the macroscale of the attractor, the value of

S_{max} should be taken at somewhat less than this value. Of course, making this parameter smaller will increase accuracy; however, the density of points on the attractor determines how small it can be, and making S_{max} too small also has the unfortunate effect of greatly increasing the computation time of the algorithm. It is instructive to do some experimentation with the effect of this variable when analyzing a data set, however, we have found that the 1% rule is usually a good guess. The second scale variable is S_{min}, which sets the minimum distance that the algorithm will look for neighbors during replacement. The purpose of this parameter is to reduce the effects of very small levels of noise by eliminating the choice of neighbors which are so close that they are within the scale of distance that the noise defines. Since we deal with "clean" data sets throughout the discussion in this paper, the value of S_{min} was set quite low. For actual experimental data corrupted by noise, a good deal of experimentation with this variable is probably necessary, as it is difficult to estimate the effective scale that the noise will appear on. It should be noted that this parameter is only effective at reducing the effects of very small magnitudes of noise, as we have found that S_{min} can usually be not much larger than about 1% of S_{max}, or the algorithm has difficulty finding sufficient numbers of neighbors within the linear regime for replacement.

The last input parameter to the algorithm is T_E, which gives the evolution time (in time series steps) that a given pair of neighbors are allowed to evolve before replacement. The value of this variable can greatly effect the accuracy of the calculation of the Lyapunov exponent, for a number of reasons. If the evolution time is too short, the difference vector between the two neighboring trajectories may not have sufficient time to evolve with the dynamics on the attractor, and the frequent replacement process can introduce considerable inaccuracies. If the evolution time is too long, the neighboring points can often evolve to distances which are greater than the linearized regime, and so these contributions are also inaccurate. Additionally, for attractors which may have a multilobed structure, such as the Lorenz attractor, enormous errors can be introduced if the evolution time is sufficiently long to allow two neighboring points to eventually evolve along the two separate lobes.

To choose T_E for a time series produced by a map, one or two iterations of the map is usually a good value, as was the case for the Hénon system. For a flow, some experimentation must be done. A good general rule is that the evolution time for a flow should be on the order of $\frac{1}{2}$ to $1\frac{1}{2}$ orbital periods on the attractor, although this again can depend on the magnitude of the Lyapunov exponents. When one only has a time series to work with, an orbital period for the system can be determined by taking a power spectrum of the time series and picking the dominant feature, if any. Once a rough estimate of what the evolution time should be is obtained, it is strongly advised to calculate the Lyapunov value for a range of evolution times around the rough value. The computed values of the Lyapunov exponent versus the evolution time will usually remain flat for some range of the evolution times, and a value within this stable range is usually an accurate choice.

Using the above general guidelines, the Wolf code was used to determine the Lyapunov exponents of two sample systems for which the exponents are already known: the Hénon map and the Lorenz system. In both cases, all of the parameters could be chosen ahead of time with good confidence, with the exception of the evolution time (T_E). For this parameter, a series of runs with differing T_E values were done, as a check of the stability of the Lyapunov value with different evolution times, and to demonstrate how this may be done with other parameters for which good guesses are not available *a priori*.

For the Hénon map (whose dimension is known to be 1.26), we chose $d=2$, and $N=2000$, although about 1000 ($=30^d$) would probably suffice. Since the system is defined by a mapping, we choose $\tau_d=1$ (this is verified using the autocorrelation calculation, Fig. 8), and likewise $t_c=1$. Since the largest scale of the map is about 4, we chose S_{max} to be 0.25 to 0.05. Also, since the data are generated numerically, the only noise is from machine error, so we chose S_{min} to be a conservative 10^{-5}. Note that some experimentation was conducted with these values, but that the result of the calculation showed λ_1 was not greatly affected for parameter values within reasonable limits of the ones given, although the run times could be considerably affected for S_{max} too small. For the remaining parameter T_E we present a graph of the value of the largest exponent λ versus the value of T_E (Fig. 9). Note that there is a plateau in the value of λ at about 0.624 for values of the parameter T_E out to about 5, after which it drops off sharply. Note that even though the characteristic time for his map is 1, we see that λ_1 is stable to a reasonably large variation in T_E. The value we obtain for λ_1 is within about 3% of the value quoted by Wolf *et al.*

For the second example, the Lorenz system, a data set was generated by integrating the dynamical equations with a Runge-Kutta integrator, using a time step [=1 (sampling rate)] of about 0.03 sec. Since the characteris-

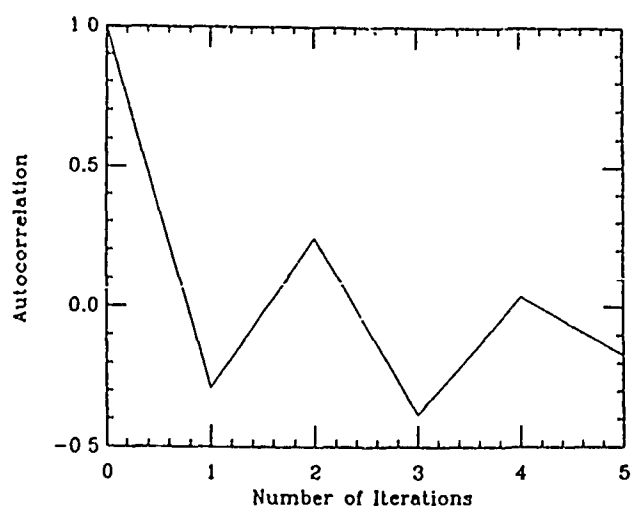

FIG. 8. Autocorrelation function for $x_1(n)$ in the Hénon system from 2000 data points.

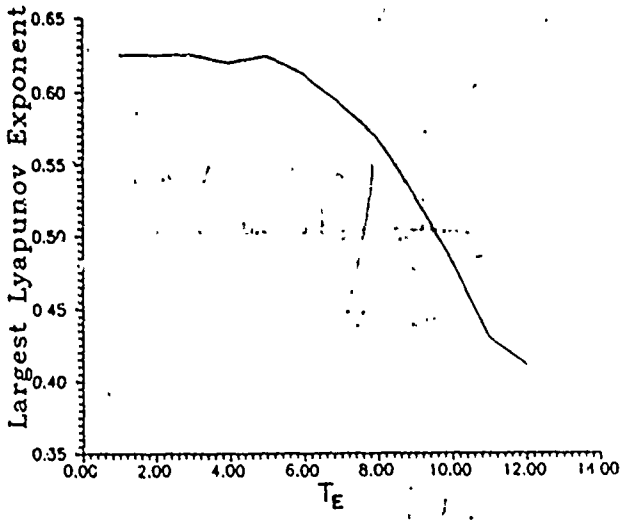

FIG. 9. λ vs T_E for the WSSV method in the Hénon map.

tic time for the Lorenz attractor is about 0.5 sec, this gives about 17 points per orbit on the attractor. The dimension of the attractor is known to be about 2.06, and an embedding dimension of $d=3$ was chosen. The minimum number of data points required, as estimated by our rule of thumb, is about 5000, so we close a set of 10 000 points ($=N$). The autocorrelation calculation (Fig. 6) suggests a value of $t_c=13$, and τ_d is 0.03. Since the maximum scale of the Lorenz attractor is about 40, $S_{\max}$ was chosen to be about 0.4 or 0.5, and $S_{\min}$ was chosen, by the same arguments as for the Hénon system, to be about 10^{-5}. As for the previous example, we calculated the largest Lyapunov exponent for a range of the last parameter T_E and these results are shown in Fig. 10. From the graph, one notes that λ_1 settles into a somewhat flat region by a value of T_E of about 16 or so (one orbital period) and remains roughly so until about 30 (two $\frac{1}{2}$ orbital periods). There is still a considerable variation in the values of λ along this region, which very likely indicates that the convergence is still not very good and a longer data set is necessary. The average value from this

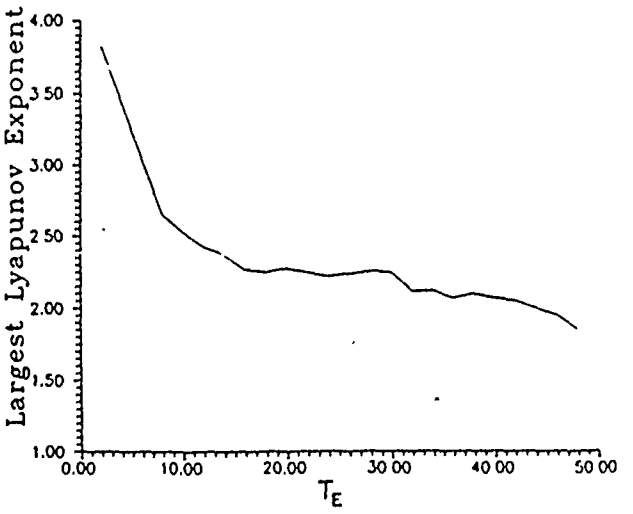

FIG. 10. λ vs T_E for the WSSV method in the Lorenz system.

regime is about 2.22, which is within 2.7% of the value 2.16 quoted by Wolf *et al.*

It is worth noting that, because of the double-lobed structure of the Lorenz attractor, the program can often be "fooled" by choosing two nearby initial orbits which wind up on different lobes of the attractor, thereby giving erroneous contributions to the averaged exponent. In this sense, the Lorenz system is a somewhat difficult case for study using the fixed-time-evolution program, and hence the results can be somewhat less accurate than would be expected.

Calculations of the largest Lyapunov exponent were carried out for other systems beside the two examples above, and in all cases the worst errors were on the order 5%, with most values being about 1–2 % of the expected exponent. We conclude that, at least for the largest positive exponent, the above code is relatively simple to use and provides reliable and reasonably accurate results. Although we have not tested them yet, more elaborate versions of the code promise greater accuracy, as well as the calculation of the rest of the positive Lyapunov exponents. The one drawback of the method is that it does not allow for calculation of the negative exponents of the spectrum, although current research suggests that it may be possible to capture at least the largest negative exponent using time reversal of the data sequence.

Some experimentation was done with calculating the largest Lyapunov exponent for a few other systems, and in all cases the worst errors were on the order of 4–5 %, with most values being within 1–2 % of the expected exponent. We can conclude from these studies that the WSSV code provides a very simple and reasonably accurate way of determining the largest Lyapunov exponent, and does not require the excessive amounts of data that some of the other algorithms seem to need. For applications where only the dominant behavior of the spreading of nearby trajectories is needed, and where it is not necessary to know the remaining Lyapunov exponents, this algorithm can prove very useful.

C. Lyapunov exponents from the map F(y,a)

Whatever method one chooses to use for determining the Lyapunov exponents from the data, and we have examined only two possible methods proposed in the literature, we must now establish a way to express these same quantities in terms of our map $\mathbf{F}(\mathbf{y},\mathbf{a})$. A direct transcription of the methods of Shimada and Nagashima,[15] Benettin, Froeschle, and Scheidecker,[16] or others would lead to a correct prescription, but not one which is easily used in the optimization or fitting we wish to do using the function $\mathbf{F}(\mathbf{y},\mathbf{a})$. The point is that one can achieve better results in this fitting if one has available a useful analytic formula for the derivatives of the constraints with respect to the parameters $\mathbf{a}$. We will choose then a slightly different way of expressing the Lyapunov exponents in terms of the map $\mathbf{F}(\mathbf{y},\mathbf{a})$ than appears in the literature. Ours may be a useful technique in itself.

Lyapunov exponents characterize the way in which neighboring points, small areas, or small volumes near the orbit of interest evolve under the mapping. To find

them one linearizes the mapping $\mathbf{y}(n+1)=\mathbf{F}(\mathbf{y}(n),\mathbf{a})$ around a given orbit $\mathbf{y}(1),\mathbf{y}(2),\ldots,\mathbf{y}(N)$. Small deviations from this orbit, call them $\delta\mathbf{y}(n)$, evolve as

$$\delta\mathbf{y}(n+1)=\mathbf{M}(\mathbf{y}(n))\delta\mathbf{y}(n) ,$$

where

$$[M(\mathbf{y})]_{kl}=\frac{\partial}{\partial y_l}F_k(\mathbf{y},\mathbf{a})$$

is evaluated along the orbit of interest. The Lyapunov exponents are found from the eigenvalues of the matrix $\mathbf{M}^K(\mathbf{y}(1))$

$$\mathbf{M}^K(\mathbf{y}(1))=\mathbf{M}(\mathbf{y}(K))\mathbf{M}(\mathbf{y}(K-1))\cdots\mathbf{M}(\mathbf{y}(1)) ,$$

which has information about the orbit generated by $\mathbf{F}(\mathbf{y},\mathbf{a})$ beginning at $\mathbf{y}(1)$. Indeed, calling the Lyapunov exponents λ_i, $i=1,2,\ldots,d$ the eigenvalues of $\mathbf{M}^K(\mathbf{y}(1))$ are $\exp(\tau K\lambda_i)$ in the limit as $K\to\infty$. The τ in this expression is the same one we set to unity in Sec. II. For the Hénon system (a map) τ is 1, while for the Lorenz system (an ODE) $\tau=0.03$ (cf. the Eckmann method in this section).

The familiar method of finding these λ_i's (Refs. 10, 15, and 16) is to apply the matrix $\mathbf{M}^K$ to an arbitrary vector $\mathbf{w}$. Then forming

$$\frac{1}{\tau K}\ln\left[\frac{\|\mathbf{M}^K\mathbf{w}\|}{\|\mathbf{w}\|}\right] \tag{9}$$

yields the largest exponent λ_1 for large K. To find the next largest exponent λ_2 one applies $\mathbf{M}^K$ to the elements of an outer product $\mathbf{w}^1\times\mathbf{w}^2$, and calculates the logarithm of the norm of this vector for large K. This gives the sum of λ_1 and λ_2. Continuing in this fashion, the full Lyapunov spectrum may, in principle, be extracted.

While the expression of the λ_i's as logarithms of the norms of various vectors to which $\mathbf{M}^K$ has been applied is correct, it presents serious problems in evaluating the derivatives with respect to the parameters $\mathbf{a}$ of the mapping $\mathbf{F}(\mathbf{y},\mathbf{a})$ from which $\mathbf{M}$ is formed. So we take a slightly different approach.

We note that the trace of the matrix $\mathbf{M}^K$ contains the information on Lyapunov exponents we desire. Our first observation is that the expression

$$\mathrm{tr}(\mathbf{M}^K)=\sum_{m=1}^{d}\exp(\tau K\lambda_m)$$

allows us to find the largest exponent λ_1 by

$$\lambda_1=\frac{1}{\tau K}\ln[\mathrm{tr}(\mathbf{M}^K)] \tag{10}$$

in the formal limit that $K\to\infty$. This expression is much more conducive to differentiation with respect to the parameters $\mathbf{a}$ (recall that $\mathbf{M}$ is a function of $\mathbf{a}$) since we have to deal with the logarithm of a simple scalar, the trace of $\mathbf{M}^K$, rather than the norm of a vector $\|\mathbf{M}^K\mathbf{w}\|$ as in Eq. (9).

One can find an expression for the next exponent λ_2 by observing that the combination

$$[\mathrm{tr}(\mathbf{M}^K)]^2-\mathrm{tr}(\mathbf{M}^{2K}) ,$$

where

$$\mathbf{M}^{2K}(\mathbf{y}(1))=\mathbf{M}(\mathbf{y}(2K))\times\mathbf{M}(\mathbf{y}(2K-1))\cdots\mathbf{M}(\mathbf{y}(2))\mathbf{M}(\mathbf{y}(1))$$

behaves as $\exp[\tau K(\lambda_1+\lambda_2)]$ for large K. So we can find the sum of $\lambda_1+\lambda_2$ by

$$\lambda_1+\lambda_2=\frac{1}{\tau K}\ln\{[\mathrm{tr}(\mathbf{M}^K)]^2-\mathrm{tr}(\mathbf{M}^{2K})\}$$

for large K. It is straightforward to construct expressions for the sum of exponents up to order m by recognizing the terms in the above logarithms as those of an expansion of the trace of the m^{th} power the matrix $(\mathbf{M}^K)_{ij}-\mathrm{tr}(M^K)\delta_{ij}$.

In any case, our procedure is now clear. Use whatever means available to evaluate the λ_i's from the data. Then form the indicated logarithms of combinations of traces of the matrices $\mathbf{M}^K$, $\mathbf{M}^{2K}$, etc. as computed from the parametrized mapping $\mathbf{F}(\mathbf{y},\mathbf{a})$. Equating the λ_i's evaluated from the data to the expressions for the λ_i's in terms of $\mathbf{F}(\mathbf{y},\mathbf{a})$ gives us a set of d constraints. We impose these constraints on our choice of the parameters $\mathbf{a}$ in conjunction with the minimization of our cost function.

Our actual practice restricts attention to the largest Lyapunov exponent λ_1 since that is the only one we know how to reliably extract from data. Thus only the trace of $\mathbf{M}^K$ is needed in our constraints. It seems to us a matter of some importance to devise accurate methods to determine the full spectrum of Lyapunov exponents from data. They would be useful in the program we are engaged in, and they act as classifiers for nonlinear dynamical systems with broadband power spectra. In the case of broadband spectra, sharp lines are not available for classifying and one must turn to the kind of dynamical invariant we have here.

IV. INVARIANT DISTRIBUTIONS ON THE ATTRACTOR

The frequency with which orbits $\mathbf{y}(n)$ visit regions of the phase space $\mathbf{R}^d$ defines an invariant distribution function, $\rho(\mathbf{y})$, which is formally defined for the mapping $\mathbf{y}(n+1)=\mathbf{F}(\mathbf{y}(n))$ as

$$\rho(\mathbf{y})=\lim_{N\to\infty}\frac{1}{N}\sum_{k=1}^{N}\delta^d(\mathbf{y}-\mathbf{F}^k(\mathbf{y}(1)))=\lim_{N\to\infty}\rho_N(\mathbf{y}) . \tag{11}$$

In a similar fashion, the invariant distribution for a numerical data set $\mathbf{y}(n)$, $n=1,\ldots,N$ is given by

$$\rho(\mathbf{y})=\lim_{N\to\infty}\frac{1}{N}\sum_{k=1}^{N}\delta^d(\mathbf{y}-\mathbf{y}(k)) . \tag{12}$$

Eckmann and Ruelle[10] discuss the features of $\rho(\mathbf{y})$ at some length. In particular, they address the question of the dependence of $\rho(\mathbf{y})$ on the initial point $\mathbf{y}(1)$. They state that any two initial points in the basin of attraction will lead to the same $\rho(\mathbf{y})$. In this sense $\rho(\mathbf{y})$ is invariant. For a dynamical system with two attractors it is possible for their basins of attraction to be intertwined in a complicated way. Any uncertainty in the initial point $\mathbf{y}(1)$

due to noise, machine round off, etc., may effect our ability to say with certainty the attractor to which a particular $\mathbf{y}(1)$ will go. Also, in the absence of noise there may be particular $\mathbf{y}(1)$'s (often, but not exclusively, on a set of measure zero in $\mathbf{R}^d$) that lead to nongeneric orbits. An example of this type of nongeneric behavior would be an unstable fixed point or periodic orbit in the presence of a strange attractor. In any event we will assume that noise levels are low and the only nongeneric orbits are unstable and of measure zero in the phase space. In this case, once a particular $\mathbf{y}(1)$ has moved beyond its transient stage the frequency with which it visits various portions of the attractor is, by definition, $\rho(\mathbf{y})$.

The complete invariant density $\rho(\mathbf{y})$ has too much information in it for our purposes. (We could not constrain a cost function to reproduce every point on the invariant density without an inordinate amount of work.) Any finite sequence of N points has a finite resolution on the attractor. That resolution is approximately N^{-1/d_A}, which is the order of the mean distance of N points on a d_A-dimensional set. Furthermore, we will never actually resolve the detailed δ-function resolution implied by Eqs. (11) and (12).

To handle this matter of finite resolution we introduce a complete orthonormal set of functions $\psi_\mu(\mathbf{y})$ defined on $\mathbf{R}^d$ which can serve as a basis set. We then expand $\rho(\mathbf{y})$ in terms of this basis

$$\rho(\mathbf{y})=\sum_{\mu=1}^{G} B_\mu \psi_\mu(\mathbf{y}) . \tag{13}$$

Truncating this expansion at some finite order ($\mu=G$) provides a finite-resolution representation corresponding to whatever information we have on $\rho(\mathbf{y})$. The coefficients B_μ will be the invariants of the dynamical process which characterize $\rho(\mathbf{y})$ within a given basis $\psi_\mu(\mathbf{y})$. After our discussion of how to select the $\psi_\mu(\mathbf{y})$'s we will establish how one extracts B_μ's both from the data vectors $\mathbf{y}(n)$ and from the parametrized map $\mathbf{F}(\mathbf{y},\mathbf{a})$. Equating the B_μ's from the data to those from the map will be our final constraints on the cost function $C(\mathbf{X},\mathbf{a})$.

While any complete orthonormal set of functions $\psi_\mu(\mathbf{y})$ would do to determine our B_μ's, some are more appealing than others. For example, Fourier series formed by taking

$$\psi_m(\mathbf{y})=e^{i\mathbf{m}\cdot\mathbf{y}}, \quad \mathbf{m}=(m_1,m_2,\ldots,m_\alpha)$$

are formally fine. However, since the attractor is occupying only a small portion of $\mathbf{R}^d$, most of the work performed by the Fourier representation of $\rho(\mathbf{y})$ will be expended in making $\rho(\mathbf{y})$ vanish off the attractor. What we seek are orthonormal functions *concentrated on the attractor,* so all the work in the expansion of $\rho(\mathbf{y})$ is expended exhibiting structure where the attractor is located. This would also result in the need for many fewer B_μ than required for Fourier series or other familiar choices for $\psi_\mu(\mathbf{y})$.

An *optimal* choice using information in the data set is constructed as follows.[26,27] Take the total data set $\mathbf{y}(n)$, $n=1,2,\ldots,$ and divide it into two portions. The first portion (of length N) will be treated as the data we are trying to model. The second portion of the data set (of length N') will be used to construct orthonormal functions. These orthonormal functions will be the $\psi_\mu(\mathbf{y})$'s that we will use in our expansion of $\rho(\mathbf{y})$, shown in Eq. (13). To explicitly construct these functions we further divide the second portion of the data into G groups of length L ($N'=LG$). Each group is a sample of the invariant attractor. If L is large enough, each sample is a significant look at $\rho(\mathbf{y})$. Treat each of the G data sets as an independent sample of $\rho(\mathbf{y})$ and form the invariant distribution for the αth sample

$$\rho_\alpha(\mathbf{y})=\frac{1}{L}\sum_{k=1}^{L}\delta^d(\mathbf{y}-\mathbf{y}(k,\alpha)) , \tag{14}$$

with $\alpha=1,2,\ldots,G$. The data point $\mathbf{y}(k,\alpha)$ is the kth member of the αth sample. Of course, the mean density of the G samples is just the total invariant density of the data set of length N',

$$\rho(\mathbf{y})=\frac{1}{G}\sum_{\alpha=1}^{G}\rho_\alpha(\mathbf{y}) . \tag{15}$$

From the G samples $\rho_\alpha(\mathbf{y})$ we form the following *phase-space correlation function*

$$R(\mathbf{z},\mathbf{w})=\frac{1}{G}\sum_{\alpha=1}^{G}\rho_\alpha(\mathbf{z})\rho_\alpha(\mathbf{w}) . \tag{16}$$

It can be shown[22,26] that the normalized eigenfunctions of this correlations function are the optimal eigenfunctions for expansion of functions localized on the attractor. Optimal means that these eigenfunctions provide the best representation in a least-squares sense of the information in $\rho(\mathbf{y})$ when expressed as a finite series in an eigenbasis. The label α is to be treated as a sampling index from a set of independent looks at the data each of which is to be thought of as selected from a uniform statistical distribution of invariant densities. The various averages over α then appear quite natural.

The requirement that $\psi_\mu(\mathbf{y})$'s be an eigenfunction of $R(\mathbf{z},\mathbf{w})$ leads to

$$\int d^d z\, R(\mathbf{w},\mathbf{z})\psi_\mu(\mathbf{z})=\mu\psi_\mu(\mathbf{w}) . \tag{17}$$

The $\psi_\mu(\mathbf{y})$'s are normalized as follows:

$$\int d^d w\, \psi_\mu(\mathbf{w})\psi_{\mu'}(\mathbf{w})=\delta_{\mu\mu'} . \tag{18}$$

As the number of samples G becomes infinite, the set of eigenfunctions becomes complete in the usual least-squares sense. If we insert Eq. (16) into Eq. (17), we see that for finite G, $R(\mathbf{w},\mathbf{z})$ becomes a finite sum of separable kernels. It is easily seen that in this case the eigenfunctions $\psi_\mu(\mathbf{y})$ must have the form

$$\psi_\mu(\mathbf{y})=\sum_{\alpha=1}^{G} C_\alpha^\mu \rho_\alpha(\mathbf{y}) . \tag{19}$$

The eigenfunctions defined in this fashion are localized near the attractor, just as we wished. This follows directly from Eq. (19) since $\psi_\mu(\mathbf{y})$ is made of the $\rho_\alpha(\mathbf{y})$'s which vanish off the attractor.

Inserting Eqs. (16) and (19) back into Eq. (17) reduces

the eigenvalue equation to a finite matrix problem. The coefficients C_α^μ are the G vectors which are eigenvectors of the $G\times G$ matrix

$$A_{\alpha\beta}=\frac{1}{G}\int d^d z\,\rho_\alpha(\mathbf{z})\rho_\beta(\mathbf{z})\,, \tag{20}$$

i.e.,

$$\sum_{\beta=1}^{G} A_{\alpha\beta}C_\beta^\mu=\mu C_\alpha^\mu\,. \tag{21}$$

We now turn to the normalization condition Eq. (18). Inserting the representation for $\psi_\mu(\mathbf{y})$ given by Eq. (19) into Eq. (18), and using the relationship between the C_α^μ's and $A_{\alpha\beta}$ given by Eqs. (20) and (21), dictates that the vectors C_α^μ obey the following normalization condition:

$$\sum_{\alpha=1}^{G} C_\alpha^\mu C_\alpha^{\mu'}=(\mu G)^{-1}\delta_{\mu\mu'}\,. \tag{22}$$

(Incidentally, this equation also shows that all the eigenvalues μ are positive.)

Formally, the elements of $\rho_\alpha(\mathbf{y})$ are δ functions. Hence, numerically speaking, computation with them is really not possible. We choose to replace $\delta^d(\mathbf{x})$ by

$$\delta^d(\mathbf{x})\rightarrow\frac{1}{(\pi\bar{\omega})^{d/2}}e^{-|\mathbf{x}|^2/\bar{\omega}}\equiv f_{\bar{\omega}}(|\mathbf{x}|)\,,$$

which, when $\bar{\omega}$ is small, represents only a small loss of resolution in calculating $\rho_\alpha(\mathbf{y})$. $f_{\bar{\omega}}$ also has the same integral as the δ function it replaces. To this approximation

$$\rho_\alpha(\mathbf{y})=\frac{1}{L}\sum_{k=1}^{L} f_{\bar{\omega}}(|\mathbf{y}-\mathbf{y}(k,\alpha)|)\,, \tag{23}$$

and Eq. (20) becomes

$$A_{\alpha\beta}=\left[\frac{1}{\pi\bar{\omega}}\right]^{d/2}\frac{1}{GL^2}\times\sum_{k,j=1}^{L}\exp[-|\mathbf{y}(k,\alpha)-\mathbf{y}(j,\beta)|^2/\bar{\omega}]\,. \tag{24}$$

We are now in a position to calculate our optimal eigenfunctions $\psi_\mu(\mathbf{y})$ from the G data sets. Use Eq. (24) to numerically calculate the $G\times G$ matrix $A_{\alpha\beta}$. Next calculate the eigenvalues μ and eigenvectors C_α^μ of this matrix, being sure to normalize them according to Eq. (22). We can then form the eigenfunctions $\psi_\mu(\mathbf{y})$ by using the normalized C_α^μ's and the $\rho_\alpha(\mathbf{y})$'s [in the form of Eq. (23)] in Eq. (19).

In Fig. 11 we show $\rho_1(\mathbf{y})$ evaluated for the Hénon attractor from a data set $L=750$ steps in length. These data are displayed on a grid of 75 points in each coordinate direction. The other densities are qualitatively similar in that they are very spikey. However, the exact position and size of the spikes varies from one sample to the next. The $\psi_\mu(\mathbf{y})$'s look like the $\rho_\alpha(\mathbf{y})$'s except that they are allowed to be negative is some regions. This is not surprising since they are composed of the $\rho_\alpha(\mathbf{y})$'s and the weights (given by the C_α^μ's) are not required to be positive [cf. Eq. (19)].

We now have a set of G orthonormal functions $\psi_\mu(\mathbf{y})$ extracted from G samples $\rho_\alpha(\mathbf{y})$ of the invariant distribution. We can use the orthonormality condition, Eq. (18), to project a particular B_μ out of Eq. (13),

$$B_\mu=\int d^d y\,\rho(\mathbf{y})\psi_\mu(\mathbf{y})\,. \tag{25}$$

Incidentally this shows the B_μ are invariants of the mapping since they are integrals of ψ_μ with the density $\rho(\mathbf{y})$ (cf. Sec. I). If we insert Eqs. (12) and (19) into this expression we get

$$B_\mu=\frac{1}{N}\sum_{\alpha=1}^{G}\sum_{k=1}^{N} C_\alpha^\mu\rho_\alpha(\mathbf{y}(k))=\frac{1}{NL}\sum_{k=1}^{N}\left[\sum_{\alpha=1}^{G}\sum_{j=1}^{L}\frac{C_\alpha^\mu}{(\pi\bar{\omega})^{d/2}}\times\exp[-|\mathbf{y}(j,\alpha)-\mathbf{y}(k)|^2/\bar{\omega}]\right], \tag{26}$$

where we have used Eq. (23). Equation (26) has been used to numerically calculate the B_μ's from the data.

This should make it clear how one evaluates the B_μ's from the $N'=GL$ data vectors $\mathbf{y}(k,\alpha)$ in $\mathbf{R}^d$. The B_μ's are the G numbers characterizing the invariant density $\rho(\mathbf{y})$ by its projection on the optimum basis vectors $\psi_\mu(\mathbf{y})$. Now we wish to see how to evaluate them from our parametrized mapping $\mathbf{F}(\mathbf{y},\mathbf{a})$. The G equalities between these two evaluations of B_μ form our final constraints on the minimization of the cost function $C(\mathbf{X},\mathbf{a})$.

To determine B_μ from the map $\mathbf{F}(\mathbf{y})$—we suppress the parameters $\mathbf{a}$ for a moment—we return to the definition of the invariant density as expressed by Eqs. (11) and (12). Call A_k the projection on $\psi_\mu(\mathbf{y})$ of each term in the sum in this equation:

$$A_k(\mu)=\int d^d y\,\psi_\mu(\mathbf{y})\delta^d(\mathbf{y}-\mathbf{F}^k(\mathbf{y}(1)))=\psi_\mu(\mathbf{F}^k(\mathbf{y}(1)))\,. \tag{27}$$

We interpret Eq. (27) as saying that $A_k(\mu)$ is the projection of $\delta^d(\mathbf{y}-\mathbf{F}^k(\mathbf{y}(1)))$ onto the orthonormal eigenfunctions $\psi_\mu(\mathbf{y})$. Using this interpretation we expand the δ function in terms of $\psi_\mu(\mathbf{y})$ to get

$$\delta^d(\mathbf{y}-\mathbf{F}^k(\mathbf{y}(1)))=\sum_{\mu=1}^{G} A_k(\mu)\psi_\mu(\mathbf{y})\,.$$

For large N, Eq. (11) can now be written as

$$\rho(\mathbf{y})=\sum_{\mu=1}^{G}\left[\frac{1}{N}\sum_{k=1}^{N} A_k(\mu)\right]\psi_\mu(\mathbf{y})\,.$$

Comparing this equation to Eq. (13) indicates that

$$B_\mu=\frac{1}{N}\sum_{k=1}^{N} A_k(\mu)=\frac{1}{N}\sum_{k=1}^{N}\psi_\mu(\mathbf{F}^k(\mathbf{y}(1)))\,.$$

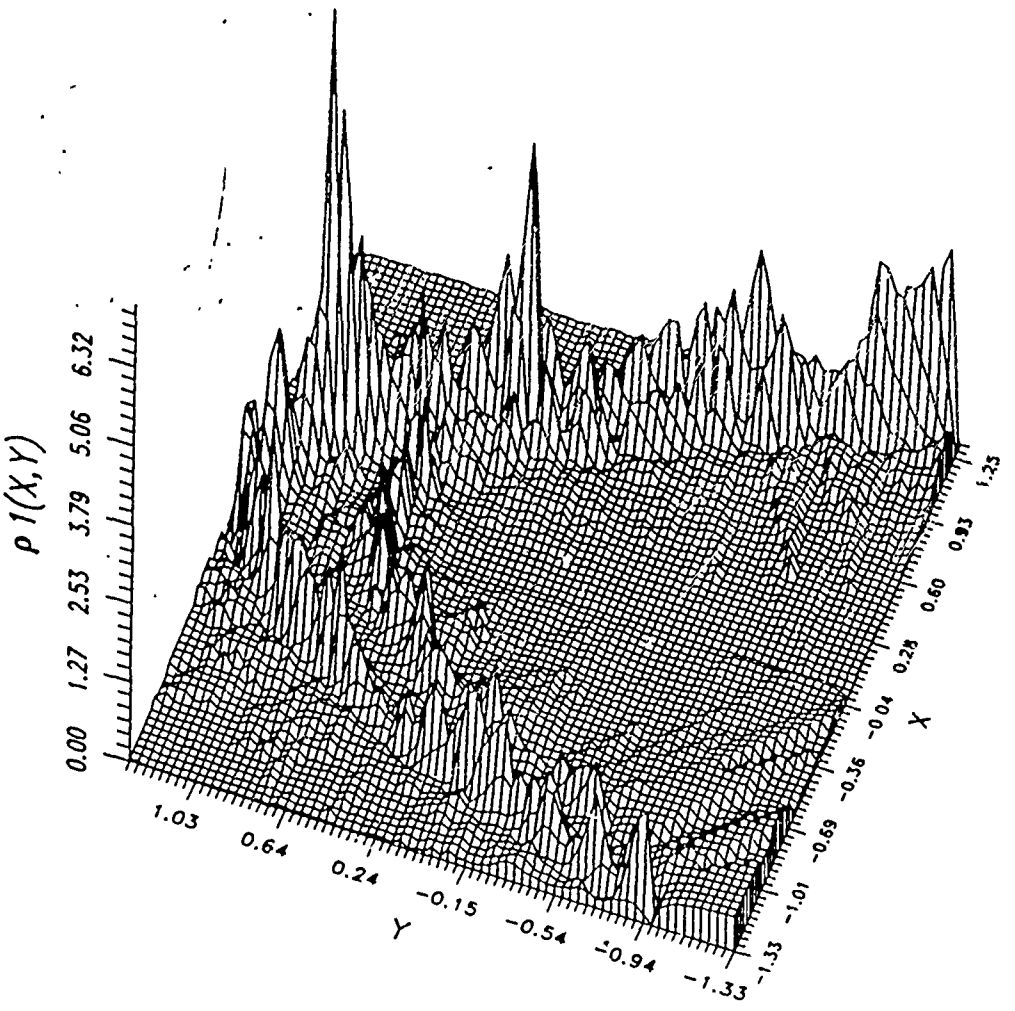

FIG. 11. First density for the Hénon map for $L=750$ and $G=5$.

From Eq. (19) and the definition of $\rho_\alpha(\mathbf{y})$ we rewrite this equation as

$$B_\mu=\frac{1}{NL}\sum_{k=1}^{N}\left[\sum_{j=1}^{L}\sum_{\alpha=1}^{G}C_\alpha^\mu\delta^d(\mathbf{y}(j,\alpha)-\mathbf{F}^k(\mathbf{y}(1)))\right].$$

This expression requires quite high powers of $\mathbf{F}(\mathbf{y},\mathbf{a})$ to be evaluated, and we cannot be confident that such high powers of the parametrized map are numerically accurate. When the map is near the correct or optimum map, then we are accurately reproducing $\mathbf{y}(k+1)$ from $\mathbf{y}(k)$ by a single application of $\mathbf{F}(\mathbf{y},\mathbf{a})$. We utilize this to replace $\mathbf{F}^k(\mathbf{y}(1))$ in the last formula with $\mathbf{F}(\mathbf{y}(k))$. The expression for B_μ's then becomes

$$\begin{aligned}B_\mu&=\frac{1}{NL}\sum_{k=1}^{N}\left[\sum_{\alpha=1}^{G}\sum_{j=1}^{L}C_\alpha^\mu\delta^d(\mathbf{y}(j,\alpha)-\mathbf{F}(\mathbf{y}(k)))\right]\\&=\frac{1}{NL}\sum_{k=1}^{N}\left[\sum_{\alpha=1}^{G}\sum_{j=1}^{L}\frac{C_\alpha^\mu}{(\pi\bar{\omega})^{d/2}}\times\exp(-|\mathbf{y}(j,\alpha)-\mathbf{F}(\mathbf{y}(k))|^2/\bar{\omega})\right],\end{aligned}\tag{28}$$

using our smoothed replacement for the δ functions in the invariant density.

This is the equation we will use to calculate the B_μ's from the map. To implement it we take all N points in the first portion of the data set and iterate them once through the map $\mathbf{F}(\mathbf{y},\mathbf{a})$. When $\mathbf{F}(\mathbf{y},\mathbf{a})$ is near the correct map iterating the data set once will result in points that are still on the attractor. We then evaluate the Gaussian, and numerically sum all the contributions between the iterates of the data set and the points in the G samples.

We close this section with an observation about the B_μ's. By combining the definition Eq. (27) with the identity

$$\delta^d(\mathbf{y}-\mathbf{F}^{k+1}(\mathbf{y}(1)))=\int d^d w\,\delta^d(\mathbf{y}-\mathbf{F}(\mathbf{w}))\delta^d(\mathbf{w}-\mathbf{F}^k(\mathbf{y}(1)))\,,$$

we can derive the recursion relation

$$A_{k+1}(\mu)=\sum_{\mu'}T_{\mu\mu'}A_k(\mu')\,,$$

in which the transition matrix T is given by

$$T_{\mu\mu'}=\int d^d y\, \psi_\mu(\mathbf{y})\psi_{\mu'}(\mathbf{F}(\mathbf{y}))\,. \tag{29}$$

In a matrix notation the recursion relation $A_{k+1}=TA_k$ leads to an expression for

$$\int d^d y\, \rho_L(\mathbf{y})\psi_\mu(\mathbf{y})=\rho_L(\mu)=\frac{1}{L}\sum_{k=1}^{L} A_k(\mu)\,,$$

which is

$$(1-T)\rho_L(\mu)=\frac{1}{L}(1-T^L)A_1\,.$$

Since $\lim_{L\to\infty}\rho_L(\mu)=\int d^d y\,\rho(\mathbf{y})\psi_\mu(\mathbf{y})=B_\mu$, this shows that the B_μ are the components of the eigenvector of T with eigenvalue unity and further that all other eigenvalues must lie within the unit circle, if the expression, Eq. (11), for $\rho(\mathbf{y})$ converges. By our assumption that the $\rho(\mathbf{y})$ we observe is unique, we infer that the eigenvalue unity of T is nondegenerate. We tried to implement this observation about the B_μ to yield a method for numerically determining them from Eq. (29) (Refs. 18 and 19), but found roundoff error undermined our efforts.

V. OPTIMIZATION OF THE CONSTRAINED COST FUNCTION: PARAMETER DETERMINATION

A. Analysis for the Hénon map

Our first application of the methods described above is to data generated by the Hénon map of the plane to itself. Data were created by iterating the map from some initial condition and discarding the first 50 points of that data set. Two data sets of $x_1(n)$ were created this way. The first had $N'=3750$ points which we divided into five groups of 750 points each. These groups were used to create the densities $\rho_\alpha(\mathbf{y})$, and the phase-space correlation function among groups was used to generate the eigenfunctions. The second data set was then used to select samples of length $N=750$, 1200, and 1752 for our analysis.

We first studied the distribution of Euclidian distances among the two vectors $\mathbf{y}(n)=(x_1(n),x_1(n+1))$, $n=1,2,\ldots,N-1$ formed from the data set. On the natural scale of the attractor, which is order unity, the minimum distance was always order 10^{-5}–10^{-4}. This led us to choose the parameter σ in our maps to be $\sigma=5\times10^{-6}$ so that each data point, at least for $N\geq500$, would have neighbors. We varied σ by a factor of 10 or so with no qualitative differences in our results. A thorough parameter search would vary σ in the constrained minimization of the cost function.

Next we chose to use four parameters $\mathbf{a}$ in our set and took the powers m_3 and m_4 in $\mathbf{F}(\mathbf{y},\mathbf{a})$ to be $m_3=4$, $m_4=5$. We did not further vary these parameters. Our choice of four $\mathbf{a}$'s rested on our knowledge that we would be constraining our cost functions by only the largest Lyapunov exponent λ_1 and the projection B_1 of $\rho(y)$ on the eigenfunction $\psi_1(\mathbf{y})$ with the largest eigenvalue. Four seemed a minimum reasonable number of parameters, and since the work required to search large parameter sets can become significant, we were content with four. In effect, we had two free parameters in $\mathbf{F}(\mathbf{y},\mathbf{a})$ when the values of λ_1 and B_1 were specified.

Our final *a priori* choice was on the values of X_j in the predictor Eq. (3). We took three terms here since we were being quite conservative in how many iterations of the map $\mathbf{F}(\mathbf{y},\mathbf{a})$ we felt we could trust. Then, further reflecting our sense that iterations of $\mathbf{F}(\mathbf{y},\mathbf{a})$ could become unreliable, we chose $X_1=0.8$, $X_2=0.1$, and $X_3=0.1$. Once again the X's could be parameters which vary in our constrained minimization. We found that varying the X's by 20% or so did not qualitatively change our results. In the case of the Lorenz attractor study discussed in Sec. V B we report results for $X_1=0.5$, $X_2=0.3$, and $X_3=0.2$, and note that the cost function changes by $\approx20\%$.

We chose to simply fix the X's for purposes of this paper. Clearly, the X's can be varied along with the $\mathbf{a}$'s, σ, and $m_3,\ldots,m_p$, if one wishes. Ours is a first try with the $\mathbf{F}(\mathbf{y},\mathbf{a})$ we have chosen in fitting the data and meeting the invariant constraints. The feasibility of accomplishing this seemed daunting enough when we set out. We expect to include many more parameters in future work in this area.

One additional important matter deserves note before we proceed to the discussion of our numerical results. The maps $\mathbf{F}(\mathbf{y},\mathbf{a})$ as we carry out our search over the parameters $\mathbf{a}$ have very little ability to reliably fit the given data for most $\mathbf{a}$. Only when we arrive near a (constrained) minimum of the cost function can we be very confident that our map is reasonable. Until the map is near the optimum map points in the data set are quite often mapped far off the attractor. For numerical stability in our search algorithms we need a method to identify orbits which are leaving the attractor for nonoptimal values of $\mathbf{a}$ and return them to the neighborhood of the attractor.

Maps of the form we have chosen have the feature that points far off the attractor, as defined by the data set itself, are mapped to $\mathbf{y}=0$. There is no reason to expect the origin of coordinates to lie on an attractor which has $d_A<d$ and is quite sparse in $\mathbf{R}^d$, but we choose to always translate our data set so one of its points is the origin. This changes nothing about the signal processing issues we address in this paper and makes our parameter searches numerically sensible.

With this translation of the origin, an orbit being generated by $\mathbf{F}(\mathbf{y},\mathbf{a})$, when $\mathbf{a}$ is not optimal, which tries to depart significantly from the attractor is sent back to $\mathbf{y}=0$, which is now on the attractor. When $\mathbf{a}$ is near its optimal values, this feature is operationally unimportant because the map is tracking the data very accurately.

Our experience indicates that if one is trying to create global maps $\mathbf{F}(\mathbf{y},\mathbf{a})$, as we are here, some form of "orbit reinjection" will be required to give numerical sense to the whole process of searching parameter space to minimize the cost function. The problem becomes more important as d grows, since the attractor of dimension $d_a<d$ occupies "less and less" of the full volume of the phase space. If one is making "fits" to the data by numerous local or nearly local polynomial maps as in the work of FS, the issue raised here is absent. Global maps

have an economy of parameters and a potential ease of interpretation; local maps appear to have an advantage of calculational speed. We have no overall judgment of a way to choose between these alternatives.

Our results for data from the Hénon map are shown in Table III. The parameter searches were carried out using the FORTRAN package NPSOL.[41] One of its authors, Gill, was kind enough to consult with us extensively on its use and on the interpretation of its output. For each value of the number of points in the data set, namely, $N=750$, 1200, and 1752, we report seven quantities for each of three cases: (1) unconstrained minimization of the cost function; (2) minimization constrained by λ_1 ($\lambda_1^{\text{map}}=0.408$); and (3) minimization constrained by both λ_1 and B_1. In each case we report the value of the cost function normalized by the sum of the squares of the Euclidian lengths over all data vectors, the values of the a's at the minimum cost function, and the deviations $\Delta\lambda_1$ and ΔB_1 from the values of λ_1^{data} and B_1^{data} determined by the data. The allowed tolerances on these deviations are set in NPSOL by the user. We typically required the relative magnitudes of $\Delta\lambda_1$ to λ_1 and the same for B_1 to be in the range 0.5–5 %. This is not a limitation of NPSOL, but it seemed quite accurate enough for our purposes.

A look at Table III reveals a consistent pattern. Unconstrained optimization resulted in a cost function with a rms deviation of our predictor from the data of 0.1% or smaller. Not surprisingly when we track the data so accurately, the value of B_1 comes out quite precise. The value of λ_1^{map} for this best least-squares fit is remarkably bad. Indeed, in our examples this quantity was actually negative, which indicates the absence of chaos for the parametric map.

When the λ_1 constraint is imposed, the parameters a change, but we regard their specific values as of incidental interest here. More important is the observation that the rms value of the cost function—the bare measure of the quality of the fit—remains about 0.5% while the Lyapunov exponent is now accurate to about 1% or better. Of course, having moved away from the very best point-to-point least-squares tracking of the data, the accuracy of B_1 degrades to $\approx 10\%$. Finally, imposing both constraints we achieve 0.5% or so in the rms error for the cost function, highly accurate λ_1, and somewhat better B_1 values.

The message of these calculations is that the procedure we outlined in this paper is both feasible and highly accurate. The few scalar numbers, the cost function, λ_1, and B_1 do not tell the whole story. One can take the map with the optimum a's and calculate a new orbit starting from some new phase space point $\mathbf{y}^{\text{new}}(1)$: $y^{\text{new}}(1)$, $y^{\text{new}}(2)$, . . . , and compare the new orbit to that generated by the Hénon map starting with the same initial point. The data so generated look the same when plotted as a sequence of two vectors, but this temporal representation contains very little useful information, so we do not show it.

What is more important is the fact that our predictor

$$\mathbf{y}(m+1)=\sum_{k=1}^{L} X_k \mathbf{F}^k(\mathbf{y}(m-k+1),\mathbf{a}) \tag{30}$$

accurately predicts. We have taken numerous points

TABLE III. Optimization results for Hénon map data. $C(\mathbf{X},\mathbf{a})$ is shown with and without invariant constraints.

$$\mathbf{F}(\mathbf{y},\mathbf{a})=\sum_{j=1}^{N-1}\mathbf{y}(j+1)g(\mathbf{y},\mathbf{y}(j);\mathbf{a})$$

$$C(\mathbf{X},\mathbf{a})=\frac{\sum_{k=L}^{N-1}|\mathbf{y}(k+1)-\sum_{j=1}^{L}x_j\mathbf{F}^j(\mathbf{y}(k-j+1),\mathbf{a})|^2}{\sum_{n=1}^{N}\mathbf{y}(n)\cdot\mathbf{y}(n)}$$

$X_1=0.8 \quad X_2=0.1 \quad X_3=0.1$

Number of points=750 $\lambda_1^{\text{data}}=0.408$ $B_1^{\text{data}}=3.4739$

	$C(\mathbf{X},\mathbf{a})$	a_1	a_2	a_3	a_4	$\Delta\lambda_1^{\text{map}}$	ΔB_1^{map}
Unconstrained	4.016×10^{-7}	7.5347	1.3289	−0.7041	0.1485	−2.0098	-5.70×10^{-3}
λ_1	4.77×10^{-6}	6.6855	20.6948	−0.1714	0.0956	1.6×10^{-4}	0.223
λ_1,B_1	2.06×10^{-5}	0.3422	1.1169	0.3766	−0.055 86	4.23×10^{-2}	0.140

Number of points=1200 $\lambda_1^{\text{data}}=0.408$ $B_1^{\text{data}}=3.388$

	$C(\mathbf{X},\mathbf{a})$	a_1	a_2	a_3	a_4	$\Delta\lambda_1^{\text{map}}$	ΔB_1^{map}
Unconstrained	3.41×10^{-6}	7.5217	2.9658	−0.3145	0.075 02	−1.676	-9.85×10^{-3}
λ_1	1.1297×10^{-5}	8.4520	26.7454	0.2686	0.011 77	-6.32×10^{-4}	0.214
λ_1,B_1	2.38×10^{-5}	6.6093	19.6341	0.083 62	−0.010 87	-1.1×10^{-4}	0.198

Number of points=1752 $\lambda_1^{\text{data}}=0.408$ $B_1^{\text{data}}=3.369$

	$C(\mathbf{X},\mathbf{a})$	a_1	a_2	a_3	a_4	$\Delta\lambda_1^{\text{map}}$	ΔB_1^{map}
Unconstrained	3.5359×10^{-7}	8.4093	6.0546	−0.1497	0.023 15	−1.211	-5.80×10^{-3}
λ_1	1.7284×10^{-5}	3.1671	9.8576	−0.1120	0.027 43	-6.512×10^{-3}	0.2605
λ_1,B_1	2.54×10^{-5}	5.8314	18.446 12	0.1832	−0.028 18	1.005×10^{-4}	0.2466

from our data set and evolved them forward by use of the predictor. We find we are able to track the actual data to the 1% level, seven to ten steps along the orbit all around the attractor. This means that iterates of our optimum map $\mathbf{F}^k(\mathbf{y},\mathbf{a})$ are accurate to $k \approx 7-10$, far beyond our original safe choice of $k=3$. The implications of this remarkable accuracy for prediction and control of nonlinear chaotic systems are transparent.

B. Prediction for the Lorenz system

We now turn to the application of our methods to the Lorenz system, defined by Eq. (8). These equations were originally motivated by an attempt to model atmospheric phenomenon using only a few degrees-of-freedom dynamical system. It was one of the first systems known to exhibit an attractor of fractal dimension, or a strange attractor, and consequently to connect this with the apparent chaotic motion of the resultant dynamics. The primary concern in modifying our previous techniques for use on the Lorenz system will be (i) the jump to a three-dimensional embedding space, which will require much longer time series to properly fill out the attractor, and (ii) the large difference in the macroscale of the two attractors, which will require the rescaling of some of the variables we have previously defined. We will first, however, give a short review of some of the characteristics of the Lorenz system.

For the parameter values $\sigma=16.0$, $r=45.92$, and $b=4.0$, the Lorenz system possesses a strange attractor which has become one of the classic examples of nonlinear science. The structure consists of two nearly flat lobes connected, roughly at a point and angled somewhat with respect to one another. Hence the local dimension of the attractor is essentially two, however, the minimum embedding space required is three. Note that the motion of the phase-space orbits for the Lorenz systems is continuous, i.e., a flow, as opposed to that of the Hénon system which is a mapping. The discretization of the Lorenz orbits after phase-space reconstruction, and the density of points along an orbit, is therefore due to the choice of a sampling rate in the measurement of the time series of data. This sampling rate therefore can be thought of as setting a time scale in the reconstructed picture of the attractor. In turn, this time scale determines the time-delay values for the method of phase-space reconstruction used in Sec. II, the evolution times for Lyapunov exponent calculations, etc. A discussion of optimal ranges of sampling rates, and the problems which occur when sampling rates are too large or small, is given at some length by Mayer-Kress.[40] In practical applications, of course, one often has no control over the data set one is presented with, although too frequent sampling can often be remedied by simply throwing away data.

To investigate the behavior of our prediction technique on a system with a somewhat larger embedding space, we chose the Lorenz system as a test case with known parameters, as was done with the Hénon system. An "experimental" time series was generated for the Lorenz equations, Eq. (8), using the parameter values listed above, by a Runge-Kutta numerical integration scheme with a fixed time step of 0.03. A data set of the x_1 variable consisting of approximately 20 000 points (after transients) was generated. We used the same time series for all of our numerical runs.

As we have stated above (cf. Sec. II) we chose an embedding dimension of three. Note that in an actual experimental situation, a more cautious choice of four would also be reasonable, although this would have increased our computational requirements by a significant amount. For the choice of the delay time constant τ a number of different choices could be made. Since the delay time reconstruction is rather weakly dependent on this constant, provided one is within certain limits, there is no unique choice for this variable. Our final choice was motivated by the desire to have the reconstructed attractor look most like the original Lorenz attractor. This results in a time delay of two time steps. For an actual case where one would have no *a priori* sense of what the attractor looks like, the methods of Sec. II are, of course, recommended. A feel for the required density of points can also be obtained by calculating the minimum nearest-neighbor distance, and perhaps the frequencies that a range of somewhat larger neighbor distances occur, and comparing this with the "macroscale" of the attractor (i.e., the maximum ranges of the coordinates of an attractor).

The embedding dimension and delay time comprise the two parameters necessary to correctly reconstruct the dynamics of the systems attractor, and hence is the first step in setting up the prediction method. We now turn to the changes necessary in the numerical algorithm when we consider the Lorenz system.

The most significant difference between the prediction models for the Hénon system and for the Lorenz system is that of the size of the time series required for the phase-space reconstruction. Because of the increase in the dimensionality of the embedding space from two to three, the number of phase-space points required to perform our procedure increases dramatically. The reasons for this is clear. Our prediction function $\mathbf{F}(\mathbf{y},\mathbf{a})$ requires that most points have a significant number of nearby neighbors, i.e., points within distances of a few $\sqrt{\sigma}$ values so that a good "mapping" of the local phase space around a particular region is obtained. Additionally, nearby neighbors are important to obtain good numerical approximations to the gradients of the objective and constraint functions. Since the number of points required to yield a given mean nearest-neighbor distance is considerably larger for a volume than for an area, the number of points required to properly fill out the attractor is much greater for a three-dimensional embedding space. In Sec. II we presented general methods for determining the number of data vectors needed for a given embedding dimension d. For our particular analysis of the Lorenz attractor reported here, we found that the minimum number of points that gave reasonable results to be about 6000. For the numerical experiments reported in Table IV we used data sets with 6000 and with 8000 points.

One final change in the numerical parameters for the prediction code is in the number of matrices that are to be multiplied together to obtain the Lyapunov exponent

TABLE IV. Optimization results for Lorenz attractor data. $C(\mathbf{X},\mathbf{a})$ is shown with and without invariant constraints.

$$\mathbf{F}(\mathbf{y},\mathbf{a})=\sum_{j=1}^{N-1}\mathbf{y}(j+1)g(\mathbf{y},\mathbf{y}(j);\mathbf{a})$$

$$C(\mathbf{X},\mathbf{a})=\frac{\sum_{k=L}^{N-1}|\mathbf{y}(k+1)-\sum_{j=1}^{L}x_j\mathbf{F}^j(\mathbf{y}(k-j+1),\mathbf{a})|^2}{\sum_{n=1}^{N}\mathbf{y}(n)\cdot\mathbf{y}(n)}$$

$X_1=0.8 \quad X_2=0.1 \quad X_3=0.1$

Number of points $=6000 \quad \lambda_1^{\text{data}}=1.51$

	$C(\mathbf{X},\mathbf{a})$	a_1	a_2	a_3	a_4	$\Delta\lambda_1^{\text{map}}$
Unconstrained	2.51672×10^{-5}	57.7977	0.08768	10.2388	−0.04358	-4.100×10^{-2}
λ_1	2.51672×10^{-5}	57.7977	0.09044	10.2388	−0.04335	-1.000×10^{-2}

Number of points $=8000 \quad \lambda_1^{\text{data}}=1.51$

	$C(\mathbf{X},\mathbf{a})$	a_1	a_2	a_3	a_4	$\Delta\lambda_1^{\text{map}}$
Unconstrained	1.87051×10^{-5}	57.7977	0.08589	10.2392	−0.03677	-6.162×10^{-2}
λ_1	1.87051×10^{-5}	57.7976	0.09352	10.2392	−0.03762	-1.101×10^{-2}

$X_1=0.5 \quad X_2=0.3 \quad X_3=0.2$

Number of points $=6000 \quad \lambda_1^{\text{data}}=1.51$

	$C(\mathbf{X},\mathbf{a})$	a_1	a_2	a_3	a_4	$\Delta\lambda_1^{\text{map}}$
Unconstrained	3.22371×10^{-5}	37.1460	0.0201	0.4224	0.00	−1.0706
λ_1	3.22372×10^{-5}	37.1459	0.0581	0.4224	0.02016	-1.000×10^{-2}

Number of points $=8000 \quad \lambda_1^{\text{data}}=1.51$

	$C(\mathbf{X},\mathbf{a})$	a_1	a_2	a_3	a_4	$\Delta\lambda_1^{\text{map}}$
Unconstrained	2.39596×10^{-5}	81.1459	0.05589	2.4222	0.00	−0.8307
λ_1	2.39597×10^{-5}	81.1456	−0.1270	2.4226	9.0809×10^{-4}	-9.999×10^{-3}

from the mapping function. Since each iteration of the Hénon map represents a significant evolution of the system, the multiplication of 500 Jacobian matrices for the Lyapunov calculation represents a good average over the phase space, and results in fairly good accuracy of the final value. However, each step of the time series for the Lorenz system represents much less evolution time for the dynamics. It was necessary to experiment with the number of matrices required to give good convergence. It was found that about 1000 matrix products gave a reasonably good convergence to the final value, but was still not excessively computationally intensive.

To complete the formulation of the prediction model for the Lorenz data, it is necessary to pick the exact form of the mapping and cost functions that are to be minimized. We first discuss the choice of the polynomial terms which multiply the exponential in the mapping function. These terms are defined, as for the Hénon analysis, with the intention of giving the exponential form in the mapping function a longer "tail" by adding multiplicative polynomial terms to it. As for the Hénon analysis, we chose to use four polynomial terms in the mapping function, and hence have four variables in the minimization fit. The first coefficient is, of course, the constant term, and the second again multiplies the linear term that expresses some dependence of the mapping function on the Lyapunov exponent. Therefore there remains to be determined the powers of the last two polynomial terms.

In choosing the values of the exponents of the remaining two polynomial terms, we recall that we wish to elongate the tail of the exponential term in the mapping function to make it feel more of the surrounding neighbors. However, we do not wish to make these exponents so large that we increase the scale well beyond that which we set by σ. After some experimentation, we chose $m=3$ and 6 as the two powers for the polynomial terms, although this is by no means the only possible choice.

The second set of parameters of the minimization procedure which need to be chosen are the X's which appear in the definition of the cost function Eq. (4). These coefficients weight the different iterates of the map $\mathbf{F}(\mathbf{y},\mathbf{a})$ and essentially determine how many iterates forward we wish the map to accurately reproduce the data. For the Hénon analysis, we chose three X's with values (0.8,0.1,0.1). Our choice indicates a desire to weight the first forward iterate very heavily, while giving the second and third iterates only minimal importance. This set of values was chosen primarily because the Hénon system is a mapping, and each iterate represents a large step in evolution of the original system. On the other hand, the Lorenz system produces a flow in phase space, and the time step we chose for each iterate of the time series represents a rather small amount of forward evolution of the system. Thus we choose to weight some of the multiple iterates of the map more heavily than we did for the Hénon system. We have therefore presented data for the Lorenz system with two different sets of values for these

parameters. In one case we used the original weights of the Hénon system (0.8,0.1,0.1). For the other case, we weighted the multiple iterates more heavily, namely, (0.5,0.3,0.2). Note that we could have easily chosen to take more than two multiple iterates of the system. However, for the sake of simplicity and comparison we chose to use two as for the Hénon system. We also point note that the X's like the a's could be made variables in the minimization search; we will do that in our further work in this matter.

Finally, to determine a value of the parameter σ (which sets a characteristic scale of distance over which the mapping function is influenced by neighbors), it is necessary to experiment with different values by actually doing a number of minimization runs. One can, however, make an *a priori* guess by considering two factors. The largest value that σ can possibly have will certainly be the scale of the linear regime for the system. This is very roughly about 1% of the attractors macroscale, as mentioned previously. Hence σ should be considerably smaller than this value. Additionally, the smallest value that σ can possibly attain is given by the smallest neighbor distance of the data set, and should be at least one to two orders of magnitude larger than this value. Within this range, σ must be chosen with some experimentation. We have found that typically, the value of the $C(\mathbf{X},\mathbf{a})$ at its minima will be relatively large for larger values of σ, and decreases until a threshold in σ is crossed. For values of σ smaller than the threshold value, the minima of $C(\mathbf{X},\mathbf{a})$ becomes a great deal less, sometimes by an order of magnitude or more. We recommend that σ be chosen somewhat smaller than this threshold value, however, not too much smaller as it is still desirable to have as much of the surrounding phase space as possible contribute to the mapping of each orbital point. For our experiments on the Lorenz system we used $\sigma = 1.0 \times 10^{-4}$.

Using the parameter values stated above, a search for the minima of Eq. (4) in the parameter space **a** was conducted using the NPSOL (Ref. 41) package. Since there is no general method known for determining the absolute minimum of a function using numerical methods, one generally proceeds by finding the minima after iteration for each of a large number of initial conditions, while attempting to cover a large representation of the phase space. In practice, one will usually find a number of local minima, all of which have "basins of attraction" of varying sizes. After a number of runs, one usually will gain some intuition as to which regions of the parameter space evolve to which local minima. When some confidence is gained that a large region of the parameter space has been investigated, we label the minimum with the lowest cost function value the "absolute" minimum. Of course, generally speaking, one can never be sure that one has bound the actual global minimum.

Using the time series for the Lorenz data and the parameter values we have just described, the NPSOL routine was able to find a number of minima of the cost function $C(\mathbf{X},\mathbf{a})$. There values ranged over as much as two orders of magnitude. The lowest value of the cost function found was in the neighborhood of 1.87×10^{-5}, as indicated in Table IV. In the preliminary analysis there were three minima which had almost this same value. A more detailed analysis, however, found that after many iterations of the search routine two of these minima actually evolved into the third. Using better error tolerances in NPSOL, it was found that this point actually did have a slightly lower minima. It should be noted that even though the three minima had cost functions which agreed very closely, their resulting values for the a's were much different. This is in keeping with our observation that, for a large range of parameter values around these minima, the cost function was very "flat" with respect to the parameters, i.e., $C(X,\mathbf{a})$ varied very little over a large range of a's. This has the unfortunate effect of causing the iteration procedure to proceed very slowly, since the minima were very shallow, and a large number of iterations were required to achieve the optimal solution. One possible conclusion from this is that, if one were interested in a purely least-squares fit of the map to the data, any of the parameter sets in this range were nearly as good as the optimal solution.

After the analysis just described, we performed another changing the X's changed to (0.5,0.3,0.2). These parameter values weight the later iterates of the map more heavily, and correspond to trying to make the map predict farther into the future. We did not impose the B_μ constraints on the Lorenz system, but used this system to explore the variations on the cost function and the quality of our ability to reproduce the largest Lyapunov exponent as we changed the weights X_j in the predictor. The results of these minimization searches are also presented in Table IV; both 6000 and 8000 points on the attractor are used in our example. As can be seen, the cost function for these minima are about $\frac{1}{3}$ higher than for the previous system, and this is to be expected since the later iterates, which must be inherently less accurate, now give a much larger contribution to the cost function. In terms of relative fitting error, however, these minima are still surprisingly low. The final parameter values, although significantly different from the previous system, are still similar enough to give the same general character to the fitting function.

One noticeable difference between the two different values of X's was in the fitting of the map using the Lyapunov constraint. The iteration procedure for the (0.5,0.3,0.2) system went far more quickly than for the (0.8,0.1,0.1) system. This can probably be interpreted in terms of the fact that if later iterates of the map are weighted more heavily, then the parameters result in more sensitivity of the map to the Lyapunov constraint, which usually requires longer evolution times to manifest itself for flows.

VI. SUMMARY AND FUTURE TASKS

In this paper we have given a set of procedures which one may use to process signals $x(n)$, $n = 1,2,\ldots$, having a broadband power spectrum. Using numerically generated data from the Hénon map and from the Lorenz equations we have also demonstrated explicitly the feasibility of our procedures. Processing a signal means that from the time series $x(n)$ we do the following.

Find an integer-dimensional embedding space of time lagged d vectors

$$\mathbf{y}(n)=(x(n),x(n+\tau_1),\ldots,x(n+\tau_{d-1}))\ ,$$

which fully expose the geometric structure of the attractor on which the data evolves. The attractor has dimension d_A which may be fractional. Choosing the integer $d>2d_A+1$ is guaranteed to be sufficient for this purpose, but smaller d may often work.

Find invariants of the evolution $\mathbf{y}(1),\mathbf{y}(2),\ldots,\mathbf{y}(N)$ in $\mathbf{R}^d$—specifically, the Lyapunov exponent spectrum $\lambda_1,\lambda_2,\ldots,\lambda_d$ and selected optimum moments $B_1,B_2,\ldots,B_G$, of the invariant density $\rho(\mathbf{y})$, on the attractor.

Use these vectors $\mathbf{y}(n)$ and invariants to construct a parametrized map of $\mathbf{R}^d$ to itself $\mathbf{y}\rightarrow\mathbf{F}(\mathbf{y},\mathbf{a})$, which minimizes a certain constrained least-squares cost function based on the residual errors of a nonlinear predictor

$$\mathbf{y}(m+1)=\sum_{k=1}^{L} X_k\mathbf{F}^k(\mathbf{y}(m-k+1),\mathbf{a})\ ,$$

involving iterates $\mathbf{F}^k$ of the map.

The output of the signal processing is the map $\mathbf{F}(\mathbf{y},\mathbf{a})$—both its form and the parameters $\mathbf{a}$—and the coefficients X_j in the predictor.

A map $\mathbf{F}(\mathbf{y},\mathbf{a})$ and a predictor which give very small least-squares residuals when evaluated on the data we call *reliable*. We have explicitly demonstrated in this paper that even a reliable $\mathbf{F}(\mathbf{y},\mathbf{a})$ does not necessarily reproduce invariants such as the λ_α and the B_μ discussed by us. The reason is that a least-squares tracking of a data set $\mathbf{y}(n)$ by a map $\mathbf{y}(n+1)\approx\mathbf{F}(\mathbf{y}(n),\mathbf{a})$ does not necessarily provide a good evaluation of the local tangent space mapping $M_{ij}=\partial F_i(\mathbf{y})/\partial y_j$. A map which is reliable and also gives the correct invariants we call *representational*. Our maps are representational because we constrain the least-squares minimization by the invariants. A map which closely tracks data but does not yield the dynamical invariants misses the essential ingredients which classify or identify the dynamical system underlying the data.

Another way to state our constrained optimization procedure is that the cost function to use in determining the map should not be composed only of the square of the residuals in the predictor. It should also contain terms which measure the residuals in matching the invariants determined by the data and the same quantity determined by the maps. NPSOL and other contemporary optimization routines do essentially this by a combination of Lagrange multiplier and quadratic penalty terms added to the least-squares cost function. This point of view suggests that we should not focus on the size of $C(\mathbf{X},\mathbf{a})$ as our goodness of fit criterion but on $C(\mathbf{X},\mathbf{a})+\sum_\alpha(\Delta\lambda_\alpha)^2+\sum_\mu(\Delta B_\mu)^2$. In our Tables III and IV we have reported the values of each of these quantities separately, but the sum as noted should measure the merit of our maps.

In practice, carrying out our signal processing program raises a number of issues of importance in dynamical systems as well as in the present context. The first of these is the determination of the dimension d of the embedding space in which the phase-space reconstruction $x(n)\rightarrow\mathbf{y}(n)$ takes place. We have used the correlation function Eq. (6), but the choice of a dimension at which this stops changing is quite subjective. Establishing an objective criterion would be most useful. Perhaps one of the information theoretic criteria developed in statistics for identifying the number of degrees of freedom in a data set would provide a tool here.[42] An objective criterion for establishing the time delays τ_α would also be desirable.

Methods for determining the Lyapunov spectrum $\lambda_1,\ldots,\lambda_d$ from the data are also quite important. These are classifiers of the dynamical system and a representational map must reproduce them. This is not at all a new issue as should be clear from the discussions in Sec. III. Our own work in this area, which will be reported in detail in a subsequent paper, uses *local* maps of the form of our $\mathbf{F}(\mathbf{y},\mathbf{a})$ and fits the parameters $\mathbf{a}$ and σ to the tangent map at every time step. The local tangent map $M(\mathbf{a}(n))_{ij}$ takes groups of phase-space points in the neighborhood of the orbit point $\mathbf{y}(n)$ into groups around $\mathbf{y}(n+1)$. The dependence of $\mathbf{M}$ on $\mathbf{y}$ is sensitive to the variation of $\mathbf{M}$ over the neighborhood of phase-space points. When one has short data sets and thus sparse neighborhoods, this dependence on $\mathbf{y}$ gives a better approximation to $\mathbf{M}(\mathbf{y})$ than a local constant matrix.[43] The eigenvalues of the product of the local $\mathbf{M}$'s along the orbit yield the λ_α.

As should be clear from our discussion of the structure of the parametrized map $\mathbf{F}(\mathbf{y},\mathbf{a})$, if we remain with our general form (which we do not insist on), then properties of $g(\mathbf{y},\mathbf{y}(n);\mathbf{a})$ are what we must address. Our choice in this paper has been to use scalar products of $\mathbf{y}$ and $\mathbf{y}(n)$ in forming g. These are insensitive to directional information on the attractor. The structure of neighborhoods of phase-space points near the orbit $\mathbf{y}(n)$ is not isotropic, so much of the information in our data may be used in our present choice of g. Since we want g to provide direction sensitive weights, we might wish to build in some of the local phase-space structure on the attractor. Some of this information is contained in the correlation function among points in the neighborhood of the orbit. If an orbit point $\mathbf{y}(n)$ has N_B neighbors $\mathbf{y}^\beta(n)$ within $\sqrt{\sigma}$, the correlation function is

$$W_{ij}(n)=\frac{1}{N_B}\sum_{\beta=1}^{N_B}[y^\beta(n)-y(n)]_i[y^\beta(n)-y(n)]_i\ .$$

Following a suggestion of Fukunaga[22] we would use the local correlation matrix in our $g(\mathbf{y},\mathbf{y}(n);\mathbf{a})$ by making the replacements

$$|\mathbf{y}-\mathbf{y}(n)|^2\rightarrow\sum_{i,j=1}^{d}[y-y(n)]_i W_{ij}^{-1}(n)[y-y(n)]_j$$

and

$$\mathbf{y}(n)\cdot(\mathbf{y}-\mathbf{y}(n))\rightarrow\sum_{i,j=1}^{d}y(n)_i W_{ij}^{-1}(n)[y-y(n)]_j\ .$$

This now emphasizes directions in phase space along the attractor where the correlation is larger.

In addition to these improvements in our ability to per-

form each element of our signal processing program, the application of methods established here to laboratory and field data would be quite productive. The applications would be both to classification by dynamical invariants of observed broadband signals and to prediction on those signals. Further having a clear idea now of the geometric setting in which the signal processing takes place in time domain, we can begin exploration of these methods to control of nonlinear systems.[44]

Finally, there is the matter of noise, extrinsic noise, which contaminates our broadband signal $x(n)$. Many conventional methods for identifying signals in noise rely on the distinct spectral characteristics of the two. That tool is absent for us, and we must use alternative tactics. We do not have a contribution to this important issue which we have tested out in any quantitative way. A natural framework will be the distinct dynamical characteristics of noise and chaotic motion embodied in differing d_A (finite for a chaotic attractor and filling any dimension for noise), invariant density $\rho(\mathbf{y})$ (structured for chaotic time series and homogeneous for noise), and other similar attributes. We will report on our tested ideas in this matter in future articles.

ACKNOWLEDGMENTS

We are most appreciative for productive conversations with K. Bruckner, M. Freedman, H. Levine, J. Theiler, and B. West about the material covered in this paper. This work was supported in part by U.S. Defense Advanced Research Projects Agency (DARPA), Applied and Computational Mathematics Program, under Contract No. F 49620-87-C-0117 and in part under the DARPA–University Research Initiative (URI), Contract No. N00014-86-K-0758. Some of the computation reported in this paper was carried out at the NASA Ames Research Center's Numerical Aerodynamic Simulation Program under the auspices of the Joint Program in Nonlinear Science between the University of California and NASA-Ames. Additionally, J.B.K. wishes to acknowledge support of the U.S. Air Force Office of Scientific Research Grant No. AFOSR-89-0072.

[1]T. Kailath, *Linear Systems* (Prentice-Hall, Englewood Cliffs, NJ, 1980).
[2]A. Papoulis, *Signal Analysis* (McGraw-Hill, New York, 1977).
[3]Francis C. Moon, *Chaotic Vibrations* (Wiley, New York, 1987).
[4]J. M. T. Thompson, and H. B. Stewart, *Nonlinear Dynamics and Chaos* (Wiley, New York, 1986).
[5]N. H. Packard, J. P. Crutchfield, J. D. Farmer, and R. S. Shaw, Phys. Rev. Lett. **45**, 712 (1980).
[6]R. Mañé, in *Dynamical Systems and Turbulence*, Vol. 898 of *Lecture Notes in Mathematics*, edited by D. Rand and L. S. Young (Springer, Berlin, 1981), pp. 230–242.
[7]F. Takens, in *Dynamical Systems and Turbulence*, Ref. 6, pp. 366–381.
[8]J.-C. Roux, R. H. Simoyi, and H. L. Swinney, Physica D **8**, 257 (1984).
[9]D. S. Broomhead and G. P. King, Physica D **20**, 217 (1986).
[10]J.-P. Eckmann and D. Ruelle, Rev. Mod. Phys. **57**, 617 (1985).
[11]J. D. Farmer and John J. Sidorowich, Phys. Rev. Lett. **59**, 845 (1987); Center for Nonlinear Studies, Los Alamos National Laboratory Report No. LA-UR-88-901, 1988.
[12]A. S. Lapedes, and R. Farber, Los Alamos National Laboratory Report No. LA-UR87-2662, 1987, and Report No. LA-UR-88-418, 1988.
[13]J. P. Crutchfield and B. S. McNamara, Complex Sys. **1**, 417 (1987).
[14]J. M. Greene and J. S. Kim, Physica D **24**, 213 (1987).
[15]I. Shimada and T. Nagashima, Prog. Theoret. Phys. **61**, 1605 (1979).
[16]G. Benettin, C. Froeschle, and J. P. Scheidecker, Phys. Rev. A **19**, 2454 (1979).
[17]V. I. Oseledec, Tr. Mosk. Mat. Obsc. Moscow Math. Soc. **19**, 17 (1968).
[18]S. Grossmann, and S. Thomae, Z. Naturforsch. **32A**, 1353 (1977).
[19]H. D. I. Abarbanel, and P. E. Latham, Phys. Lett. **89A**, 55 (1982).
[20]J. Rice, in *Computer Science and Statistics: The Interface*, edited by James E. Gentle (North-Holland, Amsterdam, 1983).
[21]B. W. Silverman, *Density Estimation for Statistics and Data Analysis* (Chapman and Hall, London, 1986).
[22]K. Fukunaga, *Introduction to Statistical Pattern Recognition* (Academic, New York, 1972), Chap. 6.
[23]J. D. Farmer and John J. Sidorowich (private communication).
[24]We are grateful to P. H. Diamond, M. H. Freedman, and H. Levine for numerous discussions about the issue of choosing the "cost function."
[25]R. A. Hummel, Comput. Graphics Image Processing **9**, 40 (1979).
[26]A. Rosenfeld and A. C. Kak, *Digital Picture Processing*, 2nd ed. (Academic, New York, 1981).
[27]W. L. Root, Proc. IEEE **75**, 1446 (1987). Root points out that the method, usually called by the names of Karhunen and Loéve, was used as early as 1950 by U. Grenander, Ark. Mat. **1**, 195 (1950); and by D. C. Youla, IRE Trans. Inf. Theory **IT-4**, 171 (1954). Clearly it has a long history.
[28]P. E. Gill, W. Murray, and M. H. Wright, *Practical Optimization* (Academic, New York, 1981); R. Fletcher, *Practical Methods of Optimization*, 2nd ed. (Wiley, Chichester, 1987).
[29]Wayne A. Fuller, *Measurement Error Models* (Wiley, New York, 1987). We thank Eric Kostelich for pointing out this work to us. Also, see E. J. Kostelich and J. A. Yorke, Phys. Rev. A **38**, 1649 (1988).
[30]John J. Sidorowich (private communication).
[31]A. M. Fraser and H. L. Swinney, Phys. Rev. **33A**, 1134 (1986); Ph.D. dissertation, University of Texas at Austin, 1988. Also see the article in Physica D **34**, 391 (1989).
[32]F. Takens, in *Proceedings of the Dynamical Systems and Bifurcations Conference, Groningen, 1984* (Springer, Berlin, 1984).
[33]P. Grassberger and I. Procaccia, Phys. Rev. Lett. **50**, 346 (1983); Physica D **9**, 189 (1983).
[34]J. Theiler, Ph.D. dissertation, California Institute of Technology, 1987 (unpublished). We have used Theiler's code for the calculation of $C(r)$ in our examples. We are most appreciative to him for lending us this code and assisting us in its use.
[35]It is usual to extract from the correlation function $D(r)$ the power ν as an estimate of the dimension of the attractor. As

Theiler (Ref. 34) has persuasively argued, this is a problematic exercise since the ν found by $\nu=\lim_{r\to 0}\log D(r)/\log r$ may not even exist. See also the recent report by Theiler, MIT Lincoln Laboratory, 1989 (unpublished). We are thankful to Dr. Theiler for sending us a copy of this paper prior to publication.

[36]M. Hénon, Commun. Math. Phys. 50, 69 (1976).

[37]E. N. Lorenz, J. Atmos. Sci. 20, 130 (1963).

[38]J.-P. Eckmann, S. O. Kamphorst, D. Ruelle, and S. Ciliberto, Phys. Rev. 34A, 4971 (1986).

[39]A. Wolf, J. B. Swift, H. L. Swinney, and J. A. Vastano, Physica 16D, 285 (1985).

[40]Gottfried Mayer-Kress, Los Alamos National Laboratory, Report No. LA-UR-87-1030, 1987 (unpublished).

[41]Philip E. Gill, Walter Murray, Michael A. Saunders, and Margaret H. Wright, *User's Guide for NPSOL (Version 4.0): A FORTRAN Package for Nonlinear Programming*, Systems Optimization Laboratory, Stanford University Technical Report No. SOL 86-2, 1986.

[42]Timothy R. C. Read and Noel A. C. Crease, *Goodness-of-Fit Statistics for Multivariate Data* (Springer-Verlag, New York, 1988), Sec. 8.3.

[43]M. Sano and Y. Sawada, Phys. Rev. Lett. 55, 1082 (1985).

[44]B. A. Huberman and E. Lumer (unpublished).

APPENDIX E

Information Theoretic Methods for Determining Minimum Embedding Dimensions for Strange Attractors

Henry D.I. Abarbanel[1]
Marine Physical Laboratory
Scripps Institution of Oceanography
and
Department of Physics

and

James B. Kadtke
Institute for Nonlinear Science

University of California, San Diego
Mail Code R-002
La Jolla, CA 92093-0402

[1]Institute for Nonlinear Science

Abstract

We present an information theoretic method for determining the minimum embedding dimension for capturing the geometry of a strange attractor from a measured time series. Unlike other methods which require a subjective judgment about thresholds, this technique determines the embedding dimension by minimizing a function which describes the information content required to produce the best model of the signal statistics as measured. We implement the method using only the sample covariance matrix of the data as input for Gaussian models. The use of higher statistics and more complex models is possible. The method is applied to two sets of computer generated data from three dimensional ordinary differential equations and to a set of laboratory data on vortex generation in the wake of an asymmetric cylinder. We find that the embedding dimension can be reliably determined using the information theoretic technique. Further, the method requires far less data than standard methods and works well in the presence of significant amounts of noise.

Analysis of measured time series from dynamical systems with strange attractors is now routinely carried out in a phase space reconstructed from time lagged scalar data $x(n)$ to make d-dimensional vectors[1]

$$\mathbf{y}(n) = [x(n), x(n+\tau), \ldots, x(n+(d-1)\tau)]. \tag{1}$$

The choice of time lags τ will be taken up below.[2] Our primary concern in this note is the dimension d of the phase space or embedding space in which one places the data $\mathbf{y}(n)$. Results of Takens[3] and Mañé[4] indicate that if the Hausdorff dimension of the attractor is d_a, then choosing d such that it is the next integer larger than $2d_a + 1$ is always sufficient to reveal the full geometrical structure of the attractor. One would like to have as small an embedding dimension d as possible for viewing the geometry of attractors for reasons of computational efficiency, and it often occurs that choosing d smaller than this sufficient condition will work. One common criterion for choosing d is to find a quantity which depends on d but becomes independent of d when d is made *large enough*. Several such indicators of d are often used which are various choices for the $q+1$ point correlation functions discussed by Paladin and Vulpiani.[5,6] Two familiar choices of these function, usually studied in scaling regimes, are for $q = -1$ when one is computing the Hausdorff dimension itself, and for $q = 1$ when one is computing the function introduced by Grassberger and Procaccia.[7] Another well tested approach is that of examining the singular value spectrum of the sample covariance matrix formed out of d-vectors constructed from lagged observational data.[8] Identifying the largest singular values above a 'noise floor' gives some indication of the dynamical degrees of freedom in the data; i.e., those not generated by noise contamination of the data.

All these criteria are subjective in that one must judge when the various quantities cease to change as the embedding dimension is increased or when some singular values of a covariance matrix have sunk below a suspected noise level. While each is likely to lead to usable choices for the embedding dimension d, when used conservatively, they may also lead to choices of d which are larger than needed. There is nothing wrong with this, in principle, since viewing the attractor in too large an embedding dimension does not lose any information about the system. It does require an unnecessarily large amount of computation, if one is planning to use the embedding space for learning about the system or predicting future behavior of the system[9] or controlling the system. Also since noise fills any dimension, working in a larger than required embedding dimension may amplify the effects of noise on one's analysis.

In this note we present an information theoretic technique adapted from statistics[10] for identifying the minimum dimension d_E required to describe the attractor. The choice of d_E is made in an objective fashion as the minimum of a function which varies with embedding dimension. One need not determine when some function becomes independent of d to choose d_E, since locating the minimum of the data description length function does that automatically.

The matter addressed in the statistics literature is this: if one observes N data vectors $\mathbf{y}(1), \mathbf{y}(2), \ldots, \mathbf{y}(N)$ and has a family of models which characterize features of these observations, how can we choose the model which best fits the data? The formulation of this 'best description' of the data is in term of a parametrized family of joint distribution functions of

the $\mathbf{y}(l)$

$$f(\mathbf{y}(1), \mathbf{y}(2), \ldots, \mathbf{y}(N)|\mathbf{a}) = f(\mathbf{Y}|\mathbf{a}). \tag{2}$$

The data vectors $\mathbf{y}(l)$ are presumed to be contaminated by measurement errors, so this kind of probabilistic description is appropriate. The parameters $\mathbf{a}$ are chosen to minimize the information required to describe the data. In this note the family of models will describe the data in terms of the rank of its sample covariance matrix, and the 'best description' corresponds to the minimum rank required to represent the features of the observed data. As we will indicate, if one has more information about the data than just the sample covariance matrix, that can be built into the family of models as well.

The manner in which one chooses the number of parameters S in the set $\mathbf{a}$ is to form the data description length (DDL) function

$$DDL(S) = -\log f(\mathbf{Y}|\hat{\mathbf{a}}) + \frac{\log N}{2}(\text{Degrees of Freedom}), \tag{3}$$

where $\hat{\mathbf{a}}$ is the maximum likelihood estimate of the parameters $\mathbf{a}$, and "Degrees of Freedom" is the number of free parameters in the set $\mathbf{a}$.

The idea is to (1) choose a model $f(\mathbf{Y}|\mathbf{a})$ which describes known features of the data, (2) find $\hat{\mathbf{a}}$ by the usual minimization of $f(\mathbf{Y}|\mathbf{a})$ with respect to the $\mathbf{a}$, and then (3) find that S which minimizes the DDL(S). The two terms in DDL(S) represent a balance between the "best fit parameters" as determined by maximum likelihood and a representation of the information contained in those parameters. The choice of model is dictated by several considerations. The primary one is the information one has about the data set. After that the ease of manipulation of the model comes into play. The more information one uses about the data, the more complicated may be the models, and the less likely one will be able to do much with it in an analytic fashion.

The connection to the problem of choosing the embedding dimension required to expose the geometry of a strange attractor is this: we know that once we have a large enough embedding dimension, we gain no information about the attractor by viewing it in still larger dimensions. We have 'unfolded' the attractor as much as possible at the minimum d, which we call d_E. The vectors in R^d, $\mathbf{y}(l)$ will have only d_E operational degrees of freedom with the remaining $d - d_E$ components filled by noise, when seen in the correct coordinate system. By choosing the minimum of DDL(S) we are identifying the minimum number of parameters required to totally describe the dynamical parts of $\mathbf{y}(l)$. Since the number of parameters is directly related to the embedding dimension by the formula we shall give below—indeed, in our case S will be the embedding dimension itself—the minimum of DDL(S) yields precisely the value d_E we seek.

In the work reported in this note we follow the lead of Wax and Kailath[11] and take the sample covariance matrix as the only information given us about the data set. This information determines the best ellipse in d space which can describe the distribution of the data in R^d. To be precise, suppose we are working in a d dimensional space with d larger than we suspect to be the best embedding dimension, d_E. From the calculation of one of the Paladin-Vulpiani[5] correlation functions, for example, one may have some clear idea of what may be too large an embedding dimension, and then one may use that to begin with. Form

the mean vector among the $y(l)$, and then the sample covariance matrix about the mean is the $d \times d$ matrix

$$R_{ab} = \frac{1}{N}\sum_{l=1}^{N} y_a(l)y_b(l), \tag{4}$$

where the indices a, b run from 1 to d, and the mean of the $y(l)$, $\frac{1}{N}\sum_{m=1}^{N} y_a(m)$, has been subtracted from each vector $y(l)$. With only this covariance matrix, we can make many different models of the way the $y(l)$ are distributed in R^d. The simplest choice of model is to take $f(\mathrm{Y}|\mathbf{a})$ to be a Gaussian. In this Gaussian we place the parameters a in a covariance matrix $\mathrm{R}(S)$ of "effective rank" S. That is, if there were no noise in the system, $\mathrm{R}(S)$ would be a $d \times d$ matrix with rank S. With noise the other $d - S$ dimensions are filled by the noise, and the actual rank is d.

We are presuming that in making the measurements y there is error of some kind. We will take that error to be additive, Gaussian, white noise with zero mean $\mathrm{n}(l)$ at each time step: $< \mathrm{n}(l) >= 0$ and $< n_a(l)n_b(l') >= \delta_{ab}\delta_{ll'}\sigma^2$. The actual measurements $y(l)$ are related to the 'true' data $z(l)$, which are d-vectors with only S independent or dynamical components, by $y(l) = z(l) + \mathrm{n}(l)$. The covariance matrix of the measured data is of the form:

$$< \mathrm{y}^T\mathrm{y} >=< \mathrm{z}^T\mathrm{z} > +\sigma^2\mathcal{I} \tag{5}$$

where $\mathcal{I}_{ab}$ is the $d \times d$ unit matrix, and $< \mathrm{z}^T\mathrm{z} >$ has rank S.

Again following Wax and Kailath[11] this suggests the parametrization of the matrix $\mathrm{R}(S)$ as

$$\mathrm{R}(S) = \sum_{m=1}^{S}(\lambda_m - \sigma^2)|V_m >< V_m| + \sigma^2\mathcal{I}, \tag{6}$$

with the λ_m and $|V_m >$ the eigenvalues and eigenvectors of $\mathrm{R}(S)$. The parameters a for our model are then the λ_m, σ^2, and the independent components of the $|V_m >$. These eigenvectors are orthonormal, so the number of independent parameters in this kind of model is $S + 1$ for the λ_m and σ^2 and $dS - S - \frac{S(S-1)}{2}$ for the components of the eigenvectors, since the vectors are real here. In Wax and Kailath[11] the vectors were complex and the counting is slightly different. The term "Degrees of Freedom" in the DDL(S) formula in this case is $S(d - \frac{S-1}{2}) + 1$.

The Gaussian joint distribution function of this model of the $y(l)$ we take to be

$$f(\mathrm{y}(1), \mathrm{y}(2), \ldots, \mathrm{y}(N)|\mathbf{a}) = \mathcal{N}\exp -[\frac{N}{2}\,\mathrm{trace}(\mathrm{R}\,\mathrm{R}(S)^{-1})], \tag{7}$$

with $\mathcal{N}$ the normalization factor

$$\mathcal{N} = \frac{1}{[(2\pi)^d \det(\mathrm{R}(S))]^{N/2}}. \tag{8}$$

This Gaussian clearly has zero mean for each $y(l)$ and the covariance of each $y(l)$ is $\mathrm{R}(S)$. It is important to note that we are not asserting here that the $y(l)$ are *actually* distributed in R^d as a Gaussian around $y = 0$, despite the fact that this function $f(\mathrm{Y}|\mathbf{a})$ has this

property. We are instead choosing a family of models for the distribution of the $y(l)$ which are commensurate with the information, namely R and $< y(l) >= 0$, we have been given. This Gaussian is the easiest model to work with among the many one might examine. With our choice of $f(\mathrm{Y}|\mathbf{a})$ we are testing the hypothesis that the sample covariance matrix R is in fact of the form $\mathrm{R}(S)$. We have not yet explored the consequence of utilizing other model functions $f(\mathrm{Y}|\mathbf{a})$ which have the same mean and covariance of the $y(l)$, but the choice of f is not dictated by the theorems from statistics,[10,11] so we have chosen the one which makes the determination of DDL(S) as easy as possible.

If we were given additional information about the data, such as the higher order cumulants around the mean, we could, of course, use that to build families of models $f(\mathrm{Y}|\mathbf{a})$ which more closely represented the observed distribution of the $y(l)$ in R^d, namely the location of the strange attractor. We will examine this possibility in our longer paper on this subject.[12] The indication is that if one works in the Fourier space dual to the $y(l)$ in which the logarithm of the characteristic function of the joint density $f(\mathrm{Y}|\mathbf{a})$ is the sum over all cumulants, then one may be able to incorporate into one's model building realistic information about the distribution of the $y(l)$. We believe, but have not demonstrated, that the information about minimum embedding dimension that we seek is well enough determined by just the sample covariance matrix that the incorporation of higher order moment information about the attractor may not be needed. In a geometric sense the embedding dimension is a crude or coarse grained characteristic of the attractor and the sample covariance matrix, as the simplest non-trivial piece of correlation information about the attractor, certainly contains that information. The higher order moments contain this information as well as more details about the fine structure of the attractor.

The maximum likelihood values of the parameters in $\mathrm{R}(S)$ are easily found from the Gaussian $f(\mathrm{Y}|\mathbf{a})$ above to be

$$\hat{\lambda}_m = \ell_m, \quad m = 1, 2, \ldots, S \tag{9}$$

$$\hat{\sigma}^2 = \frac{1}{d-S}\sum_{i=S+1}^{d} \ell_i, \tag{10}$$

and

$$|\hat{V}_m >= |C_m > \quad m = 1, 2, \ldots, S, \tag{11}$$

where the ℓ_m and the $|C_m >$ are the eigenvalues and eigenvectors of the sample covariance matrix. The eigenvalues have been ordered so $\ell_1 \geq \ell_2 \geq \ldots \geq \ell_d$. Combining these values with the determination of the number of free parameters in $\mathrm{R}(S)$ given above, we arrive at the formula for DDL(S):

$$DDL(S) = -\log[\prod_{i=S+1}^{d} \ell_i^N] + \log[\hat{\sigma}^{2(d-S)N}] + \frac{\log N}{2}(S(d - \frac{S-1}{2}) + 1). \tag{12}$$

The minimum embedding dimension d_E is the value at which $DDL(S)$ is a minimum for fixed $d > S$ and fixed N. The consistency of this criterion is demonstrated by Wax and Kailath.[11]

In applying the DDL method to a number of standard examples from the literature on chaotic systems, we have found that the DDL function can consistently estimate the embedding dimension of a signal generated by an attractor in phase space. In order that this work properly, however, it is necessary that a few parameters related to the phase space reconstruction be chosen within proper parameter regimes. If this is not done, neither this information theoretic method nor any other is likely to work. The two primary parameters to be considered are the sampling interval, τ_s, and the level of zero mean, Gaussian noise inherent in the signal. If these parameters are not in the correct range, the technique will generally not produce the correct value for d_E. These ranges are, as it turns out, rather broad, about an order of magnitude, and there are several numerical checks one can make to insure that one is in the correct regime. The procedures we present here are partly empirical, we will present further arguments in our subsequent paper.[12]

The implementation of the DDL method in this paper relies on the knowledge of the $d \times d$ sample covariance matrix, which at each step of the summation involved in its definition is generated from d consecutive points in the time series. This framework defines two time scales: (1) the 'window' width, τ_w, which is the larger scale and must be chosen to capture most of the relevant dynamics, and (2) the interpoint sampling interval $\tau_s = \frac{\tau_w}{d}$ which must be chosen to insure convergence of the singular value spectrum of the covariance matrix for a given τ_w. In choosing these parameters we used many of the ideas in Reference (8).

For a given time series, we first look at the Fourier power spectrum, from which we extract two pieces of information. The first is the estimate of the primary orbital frequency on the attractor, which can usually be taken as the dominant peak in the power spectrum. This is usually only needed to give some idea of the time scale of the dynamics. The second is the 'bandwidth' of the signal, which we estimate by picking the highest frequency beyond which the spectrum seems to be structureless. This is usually not difficult and an approximate answer suffices. We then assume that all of the relevant dynamics is contained in the frequencies below this highest frequency, f_0. The inverse of this we take as $\tau_w = \frac{1}{2f_0}$, which should now be sufficiently large to capture the dynamics. The two is from the usual Nyquist criterion.

Our choice of τ_s is now dictated by τ_w, and by two other factors. First, we wish to choose the sampling time large enough so that the information contained in successive time series points is not redundant.[7] For example, in the limit that $\tau_s \to 0$, most attractors would appear one dimensional since points are densely packed along an orbit. On the other hand, we also require that for a given τ_w there are a sufficient number $d = \frac{\tau_w}{\tau_s}$ of singular values available in the covariance matrix so the spectrum can converge. Convergence is in the sense that there are enough singular values corresponding to 'noise' that $d - d_E$ is nonzero and the correct dimension saturates at some value. This sets an upper limit on the size of the sampling interval. We have found empirically that for the lower dimensional systems with which we have been working, a value of $d \approx 10-15$ generally determines an optimal sampling rate $\tau_s = \frac{\tau_w}{d}$ for which the singular value spectrum converges quite well. We have found that sampling rates within a factor of two of this estimate produce equivalent results.

In contrast to the standard time-delay techniques for determining the embedding space, we require *a priori* a proper sampling time τ_s for the analysis. Although this may at first

seem restrictive in terms of analyzing experimental data, there is often quite considerable freedom in picking the time scale. For example, sampling rates which are too high can always be decimated to be within the proper sampling range just indicated. Sampling rates which are too low are clearly not correctable in this way, and the fact they miss some of the dynamics has to be lived with.

In addition to choosing τ_s, it is also necessary to have a sufficient amount of noise for the DDL method to be numerically accurate. This means that if the signal is initially 'clean', one may have to add in some noise to make the procedure converge well. Although as of this note we have no rigorous argument for picking the proper regime of noise level, we have determined empirically a way of estimating the noise level which seems to consistently put us into the correct regime. We have found that for an initially noise free signal, such as one might generate on a computer, the optimal choice for the noise level is about 20 dB in signal to noise defined as

$$SNR = 10 \log_{10}[\frac{< x^2 >}{< \sigma^2 >}],$$

where the numerator in the logarithm is the rms value about the mean of the signal and the denominator is the same for the Gaussian noise. This corresponds to a signal about an order of magnitude larger than the noise in rms amplitude. This level of noise seems optimal from a numerical point of view, though we have found that SNR levels in the range of 15 to 25 dB work fine. For SNR's lower than that, however, the dimension estimates will usually be systematically underestimated.

If one is analyzing a signal which has a known level of noise, then this rule gives one an idea how to add to the signal so the SNR is about correct. If one, however, has no idea of the amount of noise contamination in the data, then one must first resort to examining the singular value spectrum of the sample covariance matrix to get a hint about it. The spectrum of the singular values gives direct information about the noise level.[8] If we renormalize the singular value spectrum by dividing each by the largest singular value, then we have found that the 20 dB SNR corresponds to the smallest singular values being about 10^{-4}. Keeping the smallest singular value in this neighborhood always results in a consistent estimate of d_E.

As with the choice of τ_s, the requirement that the noise be within the range indicated may at first seem restrictive. However, for SNR smaller than 20 dB, one only need add artificially generated noise to bring it into range. Noise levels much larger than this can generally cause incorrect estimates, but they have the systematic effect of causing d_E to be underestimated by one or two.

Additionally, we note that one very good diagnostic is to check for the convergence of the DDL method by increasing the value of d = number of singular values. If the parameters are within the proper regimes, the position of the minima of the DDL function should saturate and remain constant as d increases in value to 15, 20, 25 etc. If the position of the minimum continues to increase, this is a good indication of being in the incorrect regime for at least one of the parameters.

Finally, we would like to stress that the data requirements for the DDL are quite minimal. For the parameter regimes defined above, for example, the embedding dimension of the

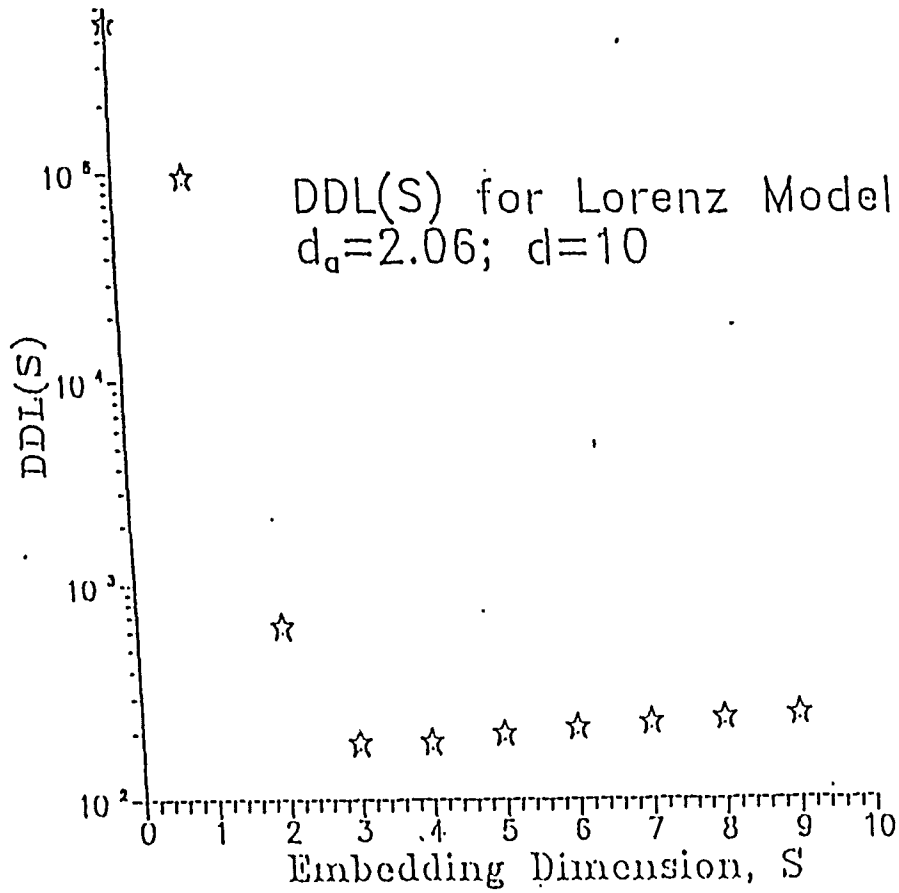

FIGURE 1

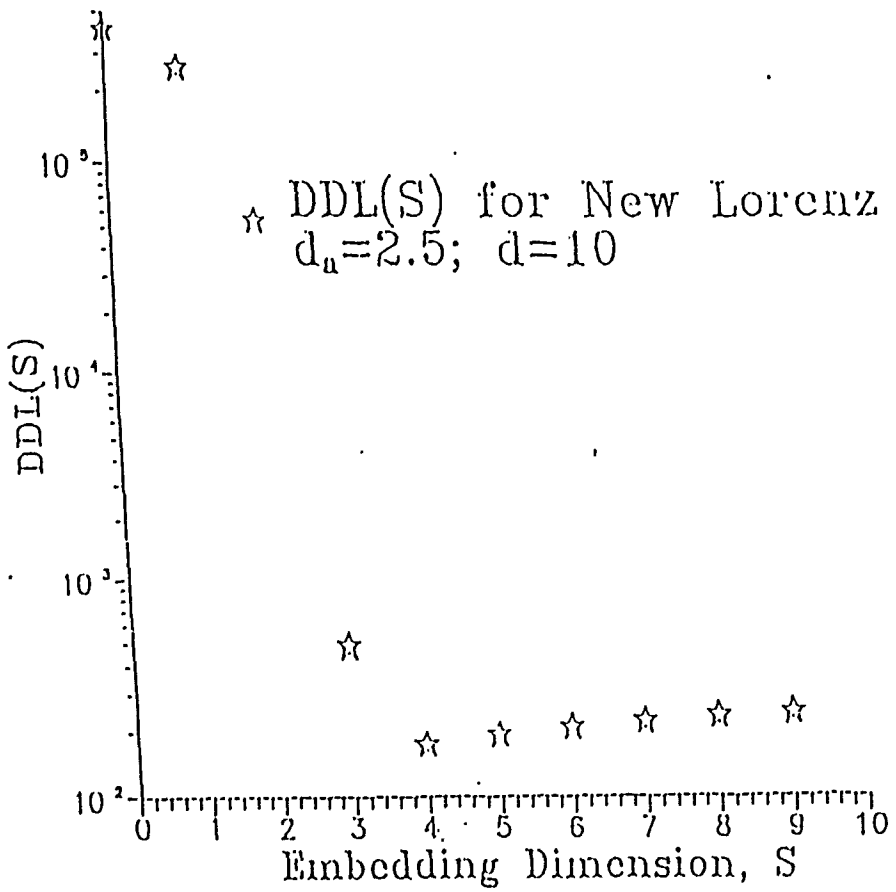

FIGURE 2

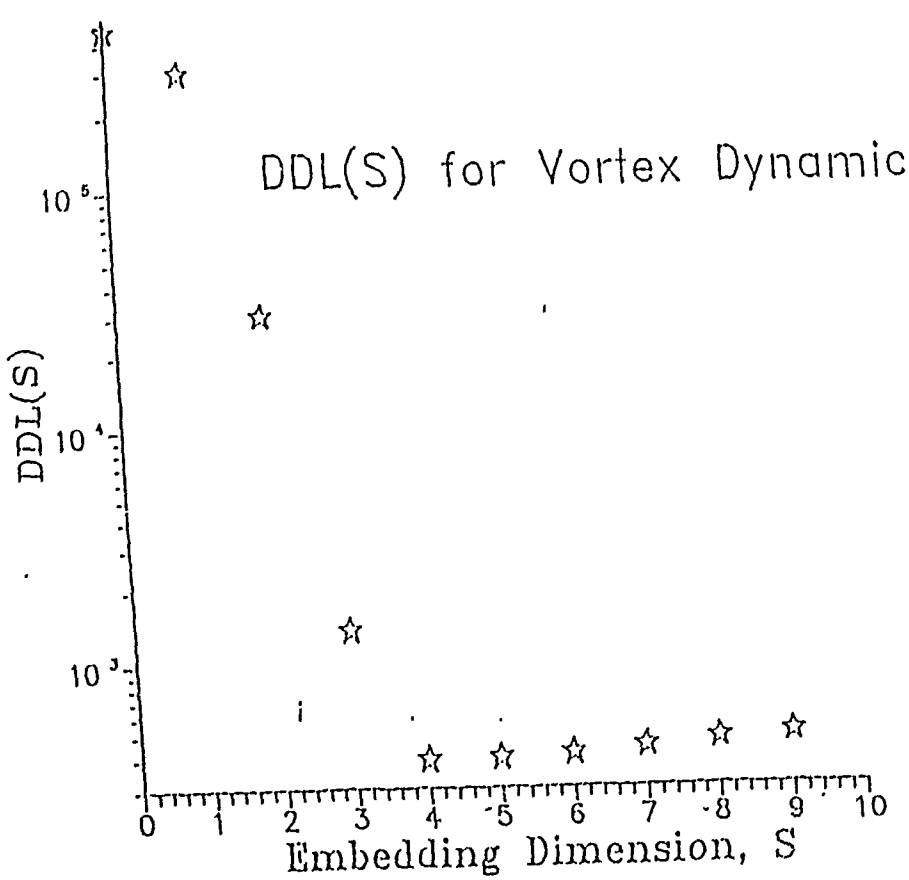
DDL(S) for Vortex Dynamics
DDL(S)
10 5
10 4
10 3
0 1 2 3 4 5 6 7 8 9 10
Embedding Dimension, S

FIGURE 3

Lorenz attractor can be estimated with as few as 2000 points in the time series. This is in marked contrast to the requirements for the correlation function method[6,7] which could require an order of magnitude more data and are far more susceptible to noise contamination.

To explore the DDL method in more detail we will discuss three examples, two of them have known embedding dimension and provide a check on the technique. The first is the classic Lorenz system,[13] which has a Hausdorff dimension of 2.06 and is known to require a three dimensional embedding space. A preliminary data set was generated from the equations using a small but arbitrary sampling time, in this case 0.03. A spectral analysis was then performed and the bandwidth limiting frequency was estimated at about 10 Hz. From this the window width τ_w was then estimated to be 0.05. We then estimate that the correct sampling rate s should be about 1/10 to 1/15 of this value, and a sampling rate of $\tau_s = 0.003$ was taken. A new data set of about 6000 points was then generated at this sampling rate, and this was used for the determination of the embedding dimension.

To choose the correct noise regime, the standard deviation about the mean of the data set was then calculated, and found to have a value of about 12.6. To obtain a noise level of 20 dB, Gaussian noise of standard deviation 1.2 was added to the signal. A check of the singular value spectrum revealed that the smallest singular values were on the order of 1×10^{-4}. The DDL function was then calculated and is shown in Figure 1. As can be seen, the minimum is obtained for an embedding dimension of three, as expected.

As a second example, we examined a second model from Lorenz[14] which represents a truncated version of a global circulation model. This system has Hausdorff dimension of about 2.5, however Lorenz has indicated that the embedding dimension necessary may be four or five[15] We again initially chose an arbitrary sampling interval of about 0.03 to generate a data set, and then performed a power spectrum analysis. From this one sees that primary period on the attractor is about 4.1, and that the band-limiting frequency was about 1.63 Hz. From this we calculate a window width $\tau_w = 0.307$, and we therefore estimate a correct sampling rate of about $\tau_s = 0.02$. A new data set of 6000 points was then generated at this sampling rate, and an estimate of the correct amount of noise was made. Since the data generated was assumed to reasonably noise-free, an amount of Gaussian noise equal to one-tenth the size of the standard deviation of the data ($< x^2 >= 0.578$) was added, or about 0.057 in magnitude. A preliminary check of the singular value spectrum revealed that the smallest singular value was on the order of 10^{-4}, so the added noise was of approximately the correct magnitude. The DDL analysis was then performed, and the results are presented in Figure 2. As can be seen, an embedding dimension of four captures the correct dynamics.

As a final example, experimental data from a vortex generation and interaction experiment in the DARPA/URI fluid dynamics laboratory[16] was analyzed. This consisted of a time series of about 50,000 data points sampled at a rate of 4000Hz. A spectral analysis reveals that the dominant period of the experiment was approximately 0.011 seconds, and the band-limiting frequency is about 188Hz. This gives a window width of approximately $\tau_w = 0.001$, and the correct sampling rate is therefore about $\tau_s = 0.0001$. However, we note that the data was sampled at 10000Hz or $\tau_s = 0.00025$, which is more than 2 times too large, and which obviously cannot be corrected. We note though that this rate is almost within the correct range, and could result in a dimension estimate slightly too large. To estimate the

correct noise regime for the system, the singular value spectrum was examined and found to be somewhat less than the range of 10^{-4}. A measurement of the standard deviation of the time series gave a value of about 295.6, which would normally indicate adding noise of a magnitude of about 30.0 to yield the 20dB. level. However, because of the amount of noise already inherent in the signal, it was found that it was necessary to add noise only of a magnitude about half that to achieve the proper regime.

The DDL method was applied to the data using these parameters, and an embedding dimension of four is found, as one can see in Figure 3. This is consistent with the dimension one would estimate from more conventional criteria,[6,7] and with the simplest models of the phenomena made by the group doing the experiment.[16] A more detailed analysis of this data will appear in a longer paper.[17]

We have thus identified an objective method for determining the minimum embedding dimension d_E for the phase space reconstruction of a strange attractor. This is quite important from the point of view of analyzing actual physical systems even though mathematically one can use and $d > d_E$ for the reconstruction. We demonstrated the simplest implementation of the *data description length* (DDL) method using only information about the sample covariance matrix of the data and showed how it works on two model examples and on experimental data taken in an experiment on vortex generation in a laboratory flow. It is very important to note that the amount of data required to determine d_E from the minimum of the DDL(S) function is very small compared to the requirements on the usual correlation functions. Also important is the fact that the method succeeds in the presence of significant levels of noise—levels at which other methods would fail. We used as few as 2000 points in a data set to determine DDL(S) for each of our examples. More data will generally be needed to accurately determine the Hausdorff dimension d_q, but if one does not require that information, as is true in many studies,[2] the extra work need not be performed. Both further detailed study of the laboratory dynamics and the richer structure of the DDL method when further moment information from the data is utilized will appear in our subsequent papers.

Acknowledgements

We are most appreciative for numerous conversations with R. Brown and P. Bryant about the material in this note. One of us (HDIA) first learned of the DDL method at an ONR Workshop in Nonlinear Problems in Physical Oceanography from Dr. M. E. Farrell and her colleagues at NADC. He is also grateful to Professor T. Kailath for a discussion of the application of these methods to this problem. The other of us (JBK) gratefully acknowledges useful discussions with C. Tirendi, MRI, UCSD and T. Hediger of NADC. Both of us are grateful to C. Lewis and M. Gharib for sharing with us their unpublished data from vortex dynamics experiments. This work was supported in part under the DARPA—University Research Initiative, URI Contract No. N00014-86-K-0758. J. B. Kadtke wishes to acknowledge support of AFOSR grant No. AFOSR-89-0072.

References

[1]Eckmann, J.-P. and D. Ruelle, "Ergodic Theory of Chaos and Strange Attractors", *Rev. Mod. Phys.* 57, 617-656 (1985)

[2]Fraser, A. M. and H. L. Swinney, "Independent Coordinates for Strange Attractors", *Phys. Rev.*, 33A, 1134-1140 (1986). "Information and Entropy in Strange Attractors". PhD Dissertation, University of Texas at Austin, May 1988. Also see the article "Reconstructing Attractors from Scalar Time Series: A Comparison of Singular System and Redundancy Criteria", *Physica* D34, 391-404 (1989).

[3]Takens, F. "Detecting Strange Attractors in Turbulence", in *Dynamical Systems and Turbulence, Warwick 1980*, eds. D. Rand and L. S. Young, Lecture Notes in Mathematics 898, (Springer, Berlin), pp. 366-381 (1981)

[4]Mañé, R. "On the Dimension of the Compact Invariant Sets of Certain Nonlinear Maps," in *Dynamical Systems and Turbulence, Warwick 1980*, eds. D. Rand and L. S. Young, Lecture Notes in Mathematics 898. (Springer, Berlin), pp. 230-242 (1981)

[5]Paladin, G. and A. Vulpiani, "Anomalous Scaling Laws in Multifractal Objects", *Phys. Repts.* 156, 147-225 (1987).

[6]Theiler, J. "Quantifying Chaos: Practical Estimation of the Correlation Dimension", PhD Dissertation, California Institute of Technology, June, 1987. Preprints from the UCSD Institute for Nonlinear Science, June, 1988. See also the recent review by Theiler, "Estimating Fractal Dimension", MIT Lincoln Laboratory Preprint, April, 1989. We are thankful to Dr. Theiler for sending us a prepublication copy of this paper.

[7]Grassberger, P. and I. Procaccia, "Characterization of Strange Attractors", *Phys. Rev. Lett.*, 50, 346-349, (1983); "Measuring the Strangeness of Strange Attractors", *Physica D* 9, 189-208 (1983).

[8]Broomhead, D. S. and G. P. King, "Extracting Qualitative Dynamics from Experimental Data", *Physica D* 20, 217 (1986). See also the very clear and useful review lecture by D. S. Broomhead and R. Jones, "Time-series analysis", *Proc. Roy. Soc. Lond. A* 423, 103-121 (1989.

[9]Abarbanel, H. D. I., R. Brown, and J. B. Kadtke, "Prediction in Chaotic Nonlinear Systems: Methods for Time Series with Broadband Fourier Spectra", INLS/UCSD Preprint 1014, May, 1989. To appear in *Phys. Rev. A*; " Prediction and System Identification in Chaotic Nonlinear Systems: Time Series with Broadband Spectra", *Phys. Lett. A* 138, 401-408 (1989).

[10]Schwartz, G., "Estimating the Dimension of a Model", *The Annals of Statistics* 6, 461-464 (1978); J. Rissanen, "Modeling by Shortest Data Description", *Automatica* 14, 465-471 (1978). See the comments by M. Stone, "Comments on model selection criteria of Akaike and Schwartz", *J. R. Statist. Soc. B* 41, 276-278 (1979).

[11]Wax, M. and T. Kailath "Detection of Signals by Information Theoretic Criteria", *IEEE Trans. Acoust., Speech, Signal Processing* ASSP-30, 387-392 (1985). These authors use the term MDL for the function we call DDL.

[12]Abarbanel, H. D. I. and J. B. Kadtke, "Information Theoretic Criteria for Phase Space Reconstruction of Strange Attractors", INLS/UCSD, to be published.

[13]Lorenz, E. N. "Deterministic Nonperiodic Flow", *J. Atmos. Sci.* 20, 130 (1963).

[14]Lorenz, E. N., "Irregularity: a fundamental property of the atmosphere", *Tellus* 36A, 98-110 (1984)

[15]Private Communication from E. N. Lorenz. February, 1989.

[16]Lewis, C. and M. Gharib, Private Communication, August, 1989.

[17]Lewis, C., M. Gharib, H. D. I. Abarbanel, and J. B. Kadtke, "Information Theoretic Dimension Analysis of Laboratory Vortex Dynamics", UCSD/INLS, to appear.